Learning from Others' Outcomes*

Alexander Wolitzky

MIT

June 7, 2017

Abstract

I develop a simple model of social learning in which players observe others' outcomes but not their actions. There is a continuum of players, and each player chooses once-and-for-all between a safe action (which succeeds with known probability) and a risky action (which succeeds with fixed but unknown probability, depending on the state of the world). The actions also differ in their costs. Before choosing, a player observes the outcomes of K others. There is always an equilibrium in which success is more likely in the good state, and this *regularity* property holds whenever the initial generation of players is not well-informed about the state. In the case of an *outcome-improving innovation* (where the risky action may yield a higher probability of success), players take the correct action as $K \to \infty$. In the case of a *cost-saving innovation* (where the risky action involves saving a cost but accepting a lower probability of success), inefficiency persists as $K \to \infty$ in any regular equilibrium. Whether inefficiency takes the form of under-adoption or over-adoption also depends on the nature of the innovation. Convergence of the population to equilibrium may be non-monotone.

*For helpful comments, I thank Daron Acemoglu, Abhijit Banerjee, Glenn Ellison, Drew Fudenberg, Ben Golub, Harry Pei, Frank Schilbach, Juuso Toikka, and seminar participants at Columbia and MIT. I thank the NSF for financial support.

1 Introduction

The development and diffusion of new technologies is a fundamental driver of economic growth. But some kinds of technologies seem to be introduced and adopted at a higher rate than others. In particular, innovations that improve observable outcomes—such as crop yields, health outcomes, or manufacturing output—while increasing costs are often more numerous and more successful than innovations that lead to worse outcomes but more than make up for this by reducing costs.

Consider the US healthcare sector, which in recent decades has experienced both rapid technological progress and rapidly rising costs (Newhouse, 1992; Cutler, 2004; Chandra and Skinner, 2012). Nelson et al. (2009) categorize 2,128 cost-effectiveness ratios from 887 medical studies published from 2002 to 2007 according to whether the studied innovations (i) increase or reduce costs, and (ii) increase or reduce quality of care, measured in quality-adjusted life years. They find that the innovations decreased cost and increased quality in 16% of cases, increased cost and decreased quality in 9% of cases, increased both cost and quality in 72% of cases, and decreased both cost and quality in 1.6% of cases. There are thus vastly more innovations that improve quality but also increase costs than innovations that save costs but reduce quality. Furthermore, several innovations in the latter category provide cost savings that greatly outweigh the corresponding reduction in quality, on any reasonable criterion. For example, treating multivessel coronary artery disease with percutaneous coronary intervention ("angioplasty") rather than coronary artery bypass graft surgery ("bypass") is reported to lead to a loss of 13 quality-adjusted life *hours*, while saving \$4,944: a saving of over \$3,000,000 per quality-adjusted life year.

Another leading example—which I will return to throughout the paper—is the overuse of agricultural fertilizer in the developing world. In the context of maize farming in Western Kenya, Duflo, Kremer, and Robinson (2008) document a strong monotonic relationship between the amount of fertilizer used per plant and the resulting yield, with 1 tsp. of fertilizer per plant (the amount recommended by the government, and also the maximum amount they report being used by farmers) producing the highest yield. But fertilizer is not free, and the relationship between amount of fertilizer per plant and *profits* is inverse U-shaped, with a maximum at 1/2 tsp. of fertilizer per plant. Indeed, most farmers (including those who use the recommended 1 tsp. per plant) actually overuse fertilizer to the point where their net returns from fertilizer utilization are strongly negative: the authors' calculations suggest that reducing fertilizer use from 1 tsp. to 1/2 tsp. per plant would increase farmers' net income by about one quarter. Given the large potential gains, why haven't farmers learned to use less fertilizer?

This paper presents a simple model of social learning, which, among other results, provides an general rationale for the scarcity of cost-saving innovations in fields such as healthcare and agriculture in the developing world (Foster and Rosenzweig, 2010) and manufacturing (Bloom and Van Reenen, 2010; Bloom et al., 2013). I argue that, when individuals learn about the quality of new innovations by observing others' outcomes, it is hard to learn about the effectiveness of a cost-saving innovation, because it is not clear if observing good outcomes is good news or bad news about the innovation. This difficulty in learning about cost-saving innovations prevents efficient adoption of these technologies, which in turn could dis-incentivize research on cost-saving technologies.[1]

To see the issue, consider the situation of a farmer who knows that using only 1/2 tsp. of fertilizer per plant will save costs while reducing the likelihood of having a successful harvest (relative to the status quo technique of using 1 tsp. of fertilizer). She is however uncertain whether the fall in expected yield will be very small (in which case using 1/2 tsp. is cost-effective) or more significant (in which case using 1 tsp. remains optimal). This situation is broadly similar to that investigated empirically by Duflo, Kremer, Robinson, and Schilbach (see Schilbach (2015) for preliminary results), who try to encourage Kenyan maize farmers to use less fertilizer by disseminating 1/2 tsp. measuring spoons.

Suppose the farmer observes her neighbors' crop yields but not their fertilizer use, and suppose she sees that average yields have fallen by only a negligible amount since the 1/2 tsp. spoons were introduced. The farmer now faces a difficult statistical inference problem:

[1]This last point may depend on the organization of the R&D sector. Public-spirited governments and non-profits presumably want to fund innovations that will be adopted efficiently, while under some market structures private firms might prefer to fund innovations that will be inefficiently overadopted. The question of how technology adoption interacts with the R&D incentives of private firms is an interesting one but is beyond the scope of the current paper. In any case, public spending is of course a critical component of R&D in many sectors, accounting for over one third of global R&D spending on biomedical technologies (Chakma et al., 2014) and over half of global R&D spending on agriculture (Pray and Fuglie, 2015).

Are relatively high yields good news about the effectiveness of using 1/2 tsp. of fertilizer, because they indicate that those farmers who used 1/2 tsp. obtained high yields? Or are they bad news, because they indicate that other farmers have persisted in using 1 tsp.?

It is not hard to see that, in order for the innovation of using only 1/2 tsp. to be adopted efficiently, observing high average yields must constitute bad news about the effectiveness of 1/2 tsp.. This follows because, if using 1/2 tsp. is adopted when it is cost-effective and is rejected when it is cost-ineffective (in favor of the higher-cost but more successful status quo of 1 tsp.), then the resulting yields will be *higher* when 1/2 tsp. is cost-*ineffective* (as 1/2 tsp. is used if and only if it is cost-effective, and 1/2 tsp. always produces smaller yields than 1 tsp.). On the other hand, as I will show, if the fraction of farmers who initially experiment with using 1/2 tsp. is the same (or close to the same) when 1/2 tsp. is cost-effective as when it is cost-ineffective (as will be the case when this first generation of farmers does not know much about the cost-effectiveness of 1/2 tsp.), then in equilibrium higher yields are always good news about the efficacy of using 1/2 tsp.. This implies that equilibrium adoption of the 1/2 tsp. measure is inefficient, regardless of the quantity of data available to farmers. If, however, one instead considers introducing an outcome-improving, cost-increasing innovation—such as a more expensive, high-yield seed variety—then in equilibrium the innovation will be adopted efficiently when enough data is available.

More generally, the key feature of my model is that learning is *outcome-based*: players observe each other's outcomes (e.g., crop yields), but not their actions (e.g., fertilizer utilization) or payoffs (e.g., profits). This contrasts with most existing models of social learning, which assume that learning is *action-based*: actions are observed, but not outcomes. Which form of learning is more appropriate depends on the application. For example, Banerjee (1992) famously considers the example of herding on the choice of a restaurant, which certainly seems like a case where action-based learning is more reasonable: one can tell at a glance if tables are full and there is a line out the door, while it is harder to tell if diners are enjoying their meals. In general, action-based learning is likely the more appropriate model for settings where actions correspond to consumption choices and outcomes to subjective utility realizations. On the other hand, in settings where actions correspond to choosing inputs in a production process, outcome-based learning is often the more natural assumption.

Indeed, in many leading economic applications, social learning seems to be better described as outcome-based rather than action-based. Consider the canonical example of the choice of agricultural or health technologies in the developing world. In this setting, it seems natural to assume that farmers can see their neighbors' crop yields more easily than they can observe what seed varieties they planted or how much fertilizer they used; or that parents can see what diseases their neighbors' children contract more clearly than they can observe what preventative medications they took.[2] Another example is given by process innovations in manufacturing, where firms can presumably observe their competitors' output levels more accurately than their production techniques. For instance, a firm might not know if a competitor has adopted a certain lean manufacturing technique, such as just-in-time delivery of inputs by suppliers, while being able to observe how frequently the competitor stocks out of its inventory.[3]

In the model, there is a large population of players, each of whom makes a once-and-for-all choice between a safe action ("status quo") and a risky action ("innovation"). The safe action yields a good outcome ("success") with known probability, while the risky action yields a good outcome with a fixed but unknown probability, which depends on the state of the world. For example, the probability that a maize crop fails to grow in a certain type of soil might be 10% when treated with 1 tsp. of fertilizer (the status quo technology), while it may be 20% or 30% when treated with 1/2 tsp. (the innovation), depending on the (unknown) effectiveness of using less fertilizer on this soil. The safe action also has a known cost advantage or disadvantage relative to the risky action. Before making her choice, a player observes the outcomes ("crop success," "crop failure") of a random sample of K other players in the population.[4] Most of my results concern the steady states ("equilibria") of this model, where the fraction of the population taking each action in each state of the

[2]See Foster and Rosenzweig (2010) for a survey of the large literature on social learning and technology adoption in development. It is interesting to note that under-adoption of cost-saving innovations appears to be a particularly severe problem in the context of healthcare in the developing world. Besides the fertilizer example, the over-prescription of antibiotics in addition to oral rehydration therapy for diarrhea is another well-known and important case; see Das et al. (2016) for some recent evidence.

[3]Bloom and Van Reenen (2010) document striking differences in the quality of management practices across firms and countries. In the context of the Indian textile industry, Bloom et al. (2013) attribute firms' failure to adopt superior production techniques to informational barriers.

[4]The outcome is assumed to be binary, which in the fertilizer application corresponds to assuming that each crop either succeeds or completely fails. This assumption is relaxed in Section 6.1.1.

world is constant over time.

A first result is that there is always an equilibrium in which success is more likely in the good state, so observing success is good news about the state. In the fertilizer example, this says that there is always an equilibrium where observing high average yields is good news about the effectiveness of using 1/2 tsp.. I call such an equilibrium—where success is good news—*regular*. One of my main results is that there is a compelling reason to focus attention on regular equilibria. In particular, I show that the equilibrium population dynamic visits only regular points whenever the initial population shares are regular, and the initial shares will be regular whenever the initial generation of players is not well-informed about the state. The intuition is that passing from a point where success is good news to a point where success is bad news requires passing through a point where one's observations are completely uninformative, and this is impossible because when observations are uninformative the population dynamic drifts back toward the initial point. This argument is illustrated graphically in Figure 1 below.[5]

My most important results address how equilibrium efficiency varies with the nature of the innovation. There turns out to be a key distinction between the case of an *outcome-improving innovation*—where the risky action may yield a higher probability of success—and that of a *cost-saving innovation*—where the risky action always yields a lower probability of success but saves costs.[6] For example, a new high-yield variety of maize (which is more expensive for farmers to use but may yield a larger harvest with higher probability than traditional maize) is an outcome-improving innovation. In contrast, using less fertilizer, angioplasty, and just-in-time input delivery (which reduce costs but may increase the likelihood of "failure") are cost-saving innovations.

In the outcome-improving case, equilibria are efficient in the limit where observations become noiseless for any $K > 1$, and equilibria are also efficient in the $K \to \infty$ limit for any amount of noise. Thus, with an outcome-improving innovation and $K > 1$, all players make the correct choice when either individual observations become perfectly informative or the

[5] There are even more direct reasons to focus attention on regular equilibria when players' samples are small: if $K = 1$ (so that each player gets only a single observation) then the unique equilibrium is regular, and if $K = 2$ then every stable equilibrium is regular.

[6] This distinction is unrelated to the distinction between good news and bad news learning familiar from the literature on strategic experimentation.

number of observations increases.[7]

In contrast, in the cost-saving case, a substantial degree of inefficiency persists in every regular equilibrium for any value of K. This result follows almost immediately from the definition of a regular equilibrium, which states that success is more likely in the good state: it is efficient to adopt the innovation in the good state but not the bad state, and, with a cost-saving innovation, this adoption pattern would lead to a higher rate of success in the *bad* state. For instance, suppose that the cost-savings from using 1/2 tsp. of fertilizer are sufficient to justify an increase in the crop failure rate from 10% to 20%, but not to 30%. Then, if players learn the state, the resulting failure rate will be 10% in the bad state and 20% in the good state, while regularity requires a higher failure rate in the bad state.

How can players possibly fail to learn the state even as they observe more and more outcomes? The only explanation is that each individual observation must become completely uninformative in the $K \to \infty$ limit: that is, the probability of success must become exactly the same in both states. The situation is thus one of confounded learning (McLennan, 1984; Smith and Sørensen, 2000). For example, suppose again that the crop failure rate is 10% with 1 tsp. of fertilizer and either 20% or 30% with 1/2 tsp., depending on the state. Then, as $K \to \infty$, the equilibrium adoption rates of 1/2 tsp. in the bad state (x_0) and the good state (x_1) must satisfy the relationship

$$\underbrace{(1 - x_1)(.1) + x_1(.2)}_{\text{failure rate in good state}} = \underbrace{(1 - x_0)(.1) + x_0(.3)}_{\text{failure rate in bad state}},$$

or $x_1 = 2x_0$, so that observing a crop failure is completely uninformative about the effectiveness of using 1/2 tsp.. This example also illustrates the inefficiency of regular equilibria in the cost-saving case: efficiency corresponds to adoption rates $(x_0 = 0, x_1 = 1)$, so the requirement that $x_1 = 2x_0$ implies substantial inefficiency.

The main finding of this paper is thus that outcome-improving innovations are generally adopted efficiently, while cost-saving innovations are adopted quite inefficiently in any regular equilibrium. I also investigate several extensions of the model and show that this main finding

[7] More precisely, for fixed $K > 1$ equilibrium is efficient in the noiseless limit in the "pure" outcome-improving case where the safe action never succeeds, while equilibrium is efficient in the $K \to \infty$ limit for any outcome-improving innovation.

continues to hold both in a range of richer physical environments (which allow for continuous actions and heterogeneous players) and richer informational environments (which allow for additional sources of information regarding the state of the world and other players' actions).

1.1 Related Literature

This paper draws on ideas and modeling techniques from several branches of the literature on social learning and experimentation.

The paper contributes to the literature on social learning and herding following Banerjee (1992), Bikhchandani, Hirshleifer, and Welch (1992), and Smith and Sørensen (2000). See Chamley (2004) for a survey. The key difference from most of this literature is that in the current paper outcomes are observable and actions are not.

On a technical level, the closest paper in this literature is Banerjee and Fudenberg (2004). Banerjee and Fudenberg assume that actions are perfectly observed, but the models are otherwise closely related. In particular, their paper introduced the approach of studying deterministic social learning dynamics in a continuum population with small sample ("word-of-mouth") learning.[8] An important difference is that under perfect observability of actions the so-called "improvement principle" applies, which states that average welfare in the population is non-decreasing over time; this follows because each player can simply copy an earlier action. The improvement principle leads to asymptotic agreement on actions and often efficiency. The assumption of perfect monitoring thus implies a rather optimistic view of social learning. In contrast, in my model the improvement principle does not hold (as will become clear), and equilibrium features interior mixing probabilities and persistent inefficiency.

A few papers consider the possibility that players might observe both actions and outcomes in social learning models. Many of these papers focus on documenting the possibility of confounded learning: outcomes becoming uninformative about the state.[9] Confounded

[8]Smith and Sørensen (2014) consider a closely related model with a discrete population.

[9]McLennan (1984) and Kiefer (1989) present classic early examples of confounded learning. Bala and Goyal (1995) and Smith and Sørensen (2000) consider confounded learning in sequential social learning models. Piketty (1995) studies confounded learning in a continuum population social learning model in a political economy application. Easley and Kiefer (1988) and Aghion, Bolton, Harris, and Jullien (1992) provide characterizations of confounded learning in single-agent experimentation problems. See Chamley (2004, Chapter 8) for a textbook treatment.

learning arises asymptotically as $K \to \infty$ in the cost-saving case of my model. However, my main contribution is characterizing how the extent and form of equilibrium inefficiency depend on the nature of the innovation, not pointing out that confounding arises per se. Another major point of contrast with this literature is that canonical papers such as McLennan (1984), Piketty (1995), and Smith and Sørensen (2000) emphasize that confounded learning is a robust equilibrium outcome in their models, while for any finite K confounded learning *cannot* occur in equilibrium in my model: instead, the fact that the equilibrium population dynamic cannot reach states where confounding would occur is a key factor in determining the equilibrium trajectory. Moreover, the possibility of confounded learning arises in my model because the adoption rate of the innovation and the success rate conditional on adoption may not be separately identified, which differs from the existing literature that assumes that actions are observable. The underlying analysis is also quite different: McLennan and Smith and Sørensen consider stochastic models and use Markov-martingale arguments to show that the equilibrium trajectory converges with positive probability to a point where learning is confounded, while my model is deterministic (conditional on the state) and, again, confounding cannot arise in equilibrium for finite K.

Perhaps more closely related, Banerjee (1993) and Acemoglu and Wolitzky (2014) consider sequential social learning models with imperfect monitoring and limited memory. As I discuss below, it is possible to view these models as discrete-time analogues of (a special case of) the current model when the entire population turns over each period. As in my paper, a key issue in these works is that players must keep track of the probability with which each action is played in each state of the world, and must use this information to correctly interpret the observed outcomes.[10] The main analysis and results in these papers are however completely different, as they focus on learning dynamics and the possibility of cycles rather than the efficiency properties of steady states.

The single-agent decision problem underlying my model is a classic two-armed bandit problem, as in the recent literature on collective experimentation (Bolton and Harris, 1999; Keller, Rady, and Cripps, 2005; see Hörner and Skrzypacz (2016) for a survey.) This lit-

[10]The same issue arises in search and matching models with incomplete information, and indeed the presence of small sample learning makes my model a kind of search model. For a striking recent contribution to the literature on information aggregation and search, see Lauermann and Wolinsky (2016).

erature is less closely related in other respects, because the fact that players in my model act only once eliminates all dynamic strategic considerations. A couple papers do however combine elements of herding and collective experimentation. Murto and Välimäki (2011) and Wagner (2017) consider stopping games with observational learning, focusing on inefficiencies due to excess delay. More closely related, Frick and Ishii (2016) study technology adoption in a model with a continuum of agents. The main modeling differences are that in their model there is no uncertainty about the share of the population using the innovation and there is no word-of-mouth learning; instead, their (long-lived) players wait to receive a public and perfectly informative signal of the state, which arrives at a higher rate when more players adopt the innovation. In terms of results, the emphasis in their paper is on the optimal timing of adoption and the shape of the resulting technology adoption curve, while there is no timing decision in my model and most of my results concern steady states. Kremer, Mansour, and Perry (2014) and Che and Hörner (2015) take a mechanism design approach to incentivizing technology adoption in related models.

2 Model and Equilibrium Existence

This section presents the model and establishes the first main result: a regular equilibrium always exists, and any equilibrium path with a regular initial point visits only regular points. This result provides the main justification for focusing on regular equilibria when analyzing equilibrium efficiency, which is the approach taken below.

2.1 Environment

There are two states, $\Theta = \{0, 1\}$, two actions, $X = \{0, 1\}$, and two outcomes, $Y = \{0, 1\}$. The prior probability of $\theta = 1$ ("the state is good," "the new technology is effective") is $p \in (0, 1)$. There is a continuum of players, and each player chooses an action once. A player who chooses action $x = 1$ ("risky action," "innovation," "experiment") when the state is θ gets outcome $y = 1$ ("success") with probability π_θ and gets outcome $y = 0$ ("failure") with probability $1 - \pi_\theta$, with $\pi_0 < \pi_1$. A player who chooses action $x = 0$ ("safe action," "status quo") gets outcome $y = 1$ with probability χ and gets outcome $y = 0$ with probability $1 - \chi$, regardless of the state. A player's payoff is $y - cx$, where $c \in \mathbb{R}$ is the cost of taking action 1 (relative to the cost of taking action 0, which is normalized to 0). Note that c can be positive or negative.

Before making her choice, each player observes the outcomes (and *not* the actions or payoffs) of a random sample of K other players. The resulting population dynamic and corresponding equilibrium concept are defined formally below.

To focus attention on the more interesting cases, throughout the paper I impose assumptions that will imply that neither always choosing the safe action (in both states) nor always choosing the risky action is an equilibrium.

The assumption that implies that always choosing the safe action is not an equilibrium is

$$p\pi_1 + (1-p)\,\pi_0 - c > \chi. \tag{1}$$

That is, the risky action is optimal under the prior. This is equivalent to assuming that p is greater than the critical belief $p^* \in (0,1)$ at which a player is indifferent between the two actions, given by $p^*\pi_1 + (1-p^*)\,\pi_0 - c = \chi$.

The assumption that implies that always choosing the risky action is not an equilibrium is

$$\left(\frac{1-\pi_0}{1-\pi_1}\right)^K > \frac{p}{1-p}\frac{1-p^*}{p^*} > 1. \tag{2}$$

As will become clear, this says that, if a player believes that everyone else experiments in both states and all K outcomes she observes are failures ($y = 0$), this is sufficiently bad news that she prefers not to experiment herself.

Note that the combination of assumptions (1) and (2) implies that the risky action is optimal in the good state while the safe action is optimal in the bad state: $\pi_1 - c > \chi > \pi_0 - c$.

2.2 Outcome-Improving vs. Cost-Saving Innovations

The model admits quite different interpretations depending on the ordering of χ and π_1.

When $\chi < \pi_1$, the risky action may yield the good outcome with higher probability than the safe action. I call this the *outcome-improving innovation* case. Note that $\chi < \pi_1$ and assumptions (1) and (2) are consistent with c being positive or negative, so that some innovations with $c < 0$ are classified as outcome-improving. The $\chi < \pi_1$ case captures the following situations:

- The safe action is planting a traditional crop variety, and the risky action is planting a new variety that could give higher or lower average yields, depending on the underlying soil suitability.
- The safe action is using a standard medical technology, and the risky action is using a more expensive new technology that promises improved health outcomes, where it is uncertain if the degree of improvement will be sufficient to justify the expense.

I refer to the extreme case where $\chi = 0$ as the *pure outcome-improving innovation* case. An example would be if the risky action corresponds to planting a completely new crop for the first time: depending on the soil conditions, the new crop may be more or less likely to grow if planted, but it will never grow without being planted.

When instead $\chi > \pi_1$, the risky action always yields the good outcome with lower probability than the safe action. Assumption (1) implies that $c < 0$ in this case: if $\pi_1 - c > \chi$ and $\chi > \pi_1$, then $c < 0$. The uncertainty facing the players is thus whether the reduced probability of receiving the good outcome is substantial enough to outweigh the lower cost of the risky action. Examples captured by this case include:

- The safe action is using the standard amount of fertilizer, and the risky action is using less fertilizer.
- The safe action is planting a traditional crop variety, and the risky action is planting an unfamiliar variety that requires less labor to cultivate.[11]
- The safe action is using traditional inventory-management techniques, and the risky action is using just-in-time delivery of inputs.
- The safe action is performing bypass, and the risky action is performing angioplasty.

[11] Bustos, Caprettini, and Ponticelli (2016) argue that genetically engineered, herbicide-resistant soybeans—which obviate the need for labor-intensive weeding—constitute a cost-saving innovation in this sense. This contrasts with the many high-yield crop varieties introduced during the Green Revolution, which are typically viewed as output-improving innovations.

- The safe action is getting a standard vaccine, and the risky action is forgoing the vaccine.

Since the $\chi > \pi_1$ case involves assuming the risk of a worse outcome to save a cost, I say that this is the case of a *cost-saving innovation.* Finally, the case where $\chi = 1$ (e.g., the standard vaccine is perfectly effective) is that of a *pure cost-saving innovation.*

2.3 Population Dynamics and Equilibrium

I introduce some notation before defining the equilibrium population dynamic.

Suppose fraction x_θ of the population takes action 1 in state θ. Then the probability that a given observation is a success in state θ equals $x_\theta \pi_\theta + (1 - x_\theta)\,\chi$. As each player's K observations are independent (conditional on θ), by Bayes' rule a player's assessment of the probability that $\theta = 1$ after observing k successes equals

$$p\,(k; x_0, x_1) = \left[1 + \frac{1-p}{p} \frac{[x_0\pi_0 + (1 - x_0)\,\chi]^k\,[1 - x_0\pi_0 - (1 - x_0)\,\chi]^{K-k}}{[x_1\pi_1 + (1 - x_1)\,\chi]^k\,[1 - x_1\pi_1 - (1 - x_1)\,\chi]^{K-k}}\right]^{-1}. \tag{3}$$

Note that $p\,(k; x_0, x_1)$ is increasing in k if and only if the probability that a given observation is a success is higher in state 1, or equivalently

$$x_1\,(\pi_1 - \chi) \geq x_0\,(\pi_0 - \chi)\,. \tag{4}$$

There are thus three possible cases (recall that p^* is the cutoff belief at which players are indifferent between the two actions):

1. (4) holds with strict inequality, there is at most one integer k^* satisfying $p\,(k^*) = p^*$, and $p\,(k) > p^*$ if and only if $k > k^*$.

2. (4) fails, there is at most one integer k^* satisfying $p\,(k^*) = p^*$, and $p\,(k) > p^*$ if and only if $k < k^*$.

3. (4) holds with equality, and $p\,(k) = p > p^*$ for all k.

I call a pair of population shares (x_0, x_1) that satisfies (4) *regular* and call population shares that satisfy the opposite of (4) *irregular*. A regular point (x_0, x_1) is thus one where observing more successes is better news about the state, in the monotone likelihood ratio sense (Milgrom, 1981). For example, in the fertilizer application, a pair of adoption rates is regular if observing crop failures is bad news about the effectiveness of using less fertilizer. Let $R \subseteq [0,1]^2$ be the set of all regular pairs (x_0, x_1).

I can now define the continuous-time equilibrium population dynamic. An equilibrium population dynamic captures the process whereby new players enter the population at rate 1 and best-respond to random samples of outcomes generated by the existing players' actions.[12] More formally, an equilibrium population dynamic specifies, at each point in time, (i) the fraction of the population experimenting in each state, (ii) a number of observed successes k^* marking the cutoff above or below which new players experiment (depending on whether or not (4) holds), and (iii) the probability with which new players experiment when they observe exactly k^* successes, if this observation leaves them indifferent.

In what follows, let $\phi_\theta(k; x)$ denote the probability that a new player observes k successes when the state is θ and fraction x of the population takes action 1:

$$\phi_\theta(k; x) = \binom{K}{k} [x\pi_\theta + (1-x)\chi]^k [1 - x\pi_\theta - (1-x)\chi]^{K-k}.$$

Definition 1 *An* equilibrium path *(or* equilibrium population dynamic*) is a list of measurable functions*

$$(x_0 : \mathbb{R}_+ \to [0,1], x_1 : \mathbb{R}_+ \to [0,1], k^* : \mathbb{R}_+ \to \{0, \dots, K\}, s : \mathbb{R}_+ \to [0,1])$$

such that

1. *Trajectories respect individual optimization: for $\theta = 0, 1$, x_θ is absolutely continuous,*

[12] There are two mathematically equivalent interpretations of the inflow of new players. The first is that new players enter at arithmetic rate 1 and existing players exit at arithmetic rate 1, keeping the overall population size constant. The second is that new players enter at exponential rate 1 and no players exit, so the overall population size increases over time. Under either interpretation, $x_\theta(t)$ will denote the fraction of the existing population taking action 1 in state θ at time t.

with derivative almost everywhere given by

$$\dot{x}_\theta(t) = \left\{ \begin{array}{c} \phi_\theta(k^*(t); x_\theta(t)) s(t) + \sum_{k=k^*(t)+1}^{K} \phi_\theta(k; x_\theta(t)) - x_\theta(t) \\ \textit{if (4) holds at } (x_0(t), x_1(t)); \\ \phi_\theta(k^*(t); x_\theta(t)) s(t) + \sum_{k=0}^{k^*(t)-1} \phi_\theta(k; x_\theta(t)) - x_\theta(t) \\ \textit{if (4) fails at } (x_0(t), x_1(t)) \end{array} \right\}. \quad (5)$$

2. *Cutoffs are consistent with Bayes' rule:*

$p(k^*(t)-1; x_0(t), x_1(t)) \leq p^* \leq p(k^*(t)+1; x_0(t), x_1(t))$ *if (4) holds at* $(x_0(t), x_1(t))$*;*

$p(k^*(t)-1; x_0(t), x_1(t)) \geq p^* \geq p(k^*(t)+1; x_0(t), x_1(t))$ *if (4) fails at* $(x_0(t), x_1(t))$*.*

3. *Decisions are optimal at the cutoff:*

$$s(t) \left\{ \begin{array}{ll} = 1 & \textit{if } p(k^*(t); x_0(t), x_1(t)) > p^* \\ = 0 & \textit{if } p(k^*(t); x_0(t), x_1(t)) < p^* \\ \in [0,1] & \textit{if } p(k^*(t); x_0(t), x_1(t)) = p^* \end{array} \right\}.$$

An equilibrium path exists for any initial point $(x_0(0), x_1(0))$.

Proposition 1 *For any point* $(\hat{x}_0, \hat{x}_1) \in [0,1]^2$*, there exists an equilibrium path* (x_0, x_1, k^*, s) *with* $(x_0(0), x_1(0)) = (\hat{x}_0, \hat{x}_1)$*.*

Proof. Define the correspondence $F : [0,1]^2 \rightrightarrows [0,1]^2$ by letting $F(x_0, x_1)$ be the set of pairs $(x'_0, x'_1) \subseteq [0,1]^2$ for which there exists $k^* \in \{0, \ldots, K\}$ and $s \in [0,1]$ such that either

$$\begin{aligned} x'_\theta &= \phi_\theta(k^*; x_\theta) s + \sum_{k=k^*+1}^{K} \phi_\theta(k; x_\theta) \text{ for } \theta = 0, 1, \\ p(k^*-1; x_0, x_1) &\leq p^* \leq p(k^*+1; x_0, x_1), \text{ and} \\ & s \left\{ \begin{array}{ll} = 1 & \text{if } p(k^*; x_0, x_1) > p^* \\ = 0 & \text{if } p(k^*; x_0, x_1) < p^* \\ \in [0,1] & \text{if } p(k^*; x_0, x_1) = p^* \end{array} \right\}, \end{aligned} \quad (6)$$

or

$$
\begin{aligned}
x'_\theta &= \phi_\theta\left(k^*; x_\theta\right) s + \sum_{k=0}^{k^*-1} \phi_\theta\left(k; x_\theta\right) \text{ for } \theta = 0, 1, \\
p\left(k^* - 1; x_0, x_1\right) &\geq p^* \geq p\left(k^* + 1; x_0, x_1\right), \text{ and (6).}
\end{aligned}
$$

Note that, if a pair of absolutely continuous functions $\left(x_0 : \mathbb{R}_+ \to [0,1], x_1 : \mathbb{R}_+ \to [0,1]\right)$ satisfies

$$
\left(\dot{x}_0(t), \dot{x}_1(t)\right) \in \left\{ \begin{array}{c} \left(\dot{x}_0, \dot{x}_1\right) : \exists \left(x'_0, x'_1\right) \in F\left(x_0(t), x_1(t)\right) \\ \text{with } \left(\dot{x}_0, \dot{x}_1\right) = \left(x'_0, x'_1\right) - \left(x_0(t), x_1(t)\right) \end{array} \right\} \tag{7}
$$

almost everywhere, then (x_0, x_1) is an equilibrium path (together with the corresponding values of k^* and s at points where (x_0, x_1) is differentiable, with the values of k^* and s at points of non-differentiability of (x_0, x_1) selected arbitrarily). I claim that F is non-empty, convex- and compact-valued, and upper hemi-continuous. The proof is routine and is deferred to Appendix B. Note also that any solution to (7) satisfies $\dot{x}_\theta(t) \in [-x_\theta(t), 1 - x_\theta(t)]$ almost everywhere, and hence cannot escape the compact set $[0, 1]$. Under these conditions, existence of a solution to the differential inclusion (7) for an arbitrary initial point is standard. (See, e.g., Aubin and Cellina (1984), Chapter 2.1, Theorem 4). ■

Most of my results concern steady states of the population dynamic.

Definition 2 *An* equilibrium *(or* steady state*) is a constant equilibrium path.*

When referring to steady states, I will abusively use the notation x_0, x_1, k^*, and s to refer to scalars rather than constant functions.

The following result simplifies the characterization of equilibria.

Proposition 2 *A vector (x_0, x_1, k^*, s) is an equilibrium if and only if either*

1. *(4) holds with strict inequality and*

$$
\begin{aligned}
x_\theta &= \phi_\theta\left(k^*; x_\theta\right) s + \sum_{k=k^*+1}^{K} \phi_\theta\left(k; x_\theta\right) \text{ for } \theta = 0, 1, \\
p\left(k^* - 1; x_0, x_1\right) &\leq p^* \leq p\left(k^* + 1; x_0, x_1\right), \text{ and (6),}
\end{aligned} \tag{8}
$$

or

2. (4) fails and

$$x_\theta = \phi_\theta(k^*; x_\theta)\, s + \sum_{k=0}^{k^*-1} \phi_\theta(k; x_\theta) \text{ for } \theta = 0, 1, \quad (9)$$

$$p(k^* - 1; x_0, x_1) \geq p^* \geq p(k^* + 1; x_0, x_1), \text{ and (6).}$$

Proof. The result follows immediately from the definitions, provided there is no equilibrium where (4) holds with equality. However, if (4) holds with equality at steady state (x_0, x_1, k^*, s) then $p(k; x_0, x_1) = p > p^*$ for all k, so the definition of a steady state implies that $x_0 = x_1 = 1$. But if $x_0 = x_1 = 1$ then (4) does not hold with equality. So there is no such equilibrium. ■

I refer to an equilibrium in which (4) holds as a *regular equilibrium* and refer to an equilibrium in which the opposite of (4) holds as an *irregular equilibrium.* Note that it is immediate that *some* (not necessarily regular) equilibrium exists. This follows because equilibria correspond to fixed points of the correspondence F introduced in the proof of Proposition 1, and F satisfies the conditions of Kakutani's fixed point theorem.

2.4 Existence and Robustness of Regular Equilibrium

To prove that a regular equilibrium exists, I define a correspondence $G_\eta : R \rightrightarrows R$, the fixed points of which correspond to regular equilibria, and show that G_η satisfies the conditions of Kakutani's fixed point theorem. The correspondence G_η generates a discrete-time version of the continuous-time population dynamic of Definition 1: for any initial pair of population shares (x_0, x_1), $G_\eta(x_0, x_1)$ is the set of population shares that can result when a small fraction η of the population dies and is replaced by new players who best-respond to samples generated by (x_0, x_1). The key idea of the theorem is that, if the initial population shares (x_0, x_1) are regular and η is sufficiently small, then every element of $G_\eta(x_0, x_1)$ is also regular: that is, the population dynamic defined by G_η can never exit the set R. The reason is that, close to the boundary between the regular and irregular regions of the unit square (i.e., close to the line L with equation $x_1(\pi_1 - \chi) - x_0(\pi_0 - \chi) = 0$), each observation is about equally likely to be a success in either state, and the players' samples therefore contain little information about the state. By (1), this implies that new players always experiment when the existing population shares are close to the boundary of R. This behavior then drives the population dynamic toward the regular point $(x_0, x_1) = (1, 1)$, and in particular drives the dynamic away from the boundary of R. Thus, as long as η is small enough so that the population dynamic cannot jump out of R without first getting close to its boundary, the dynamic can never leave R. See Figure 1 for an illustration.[13]

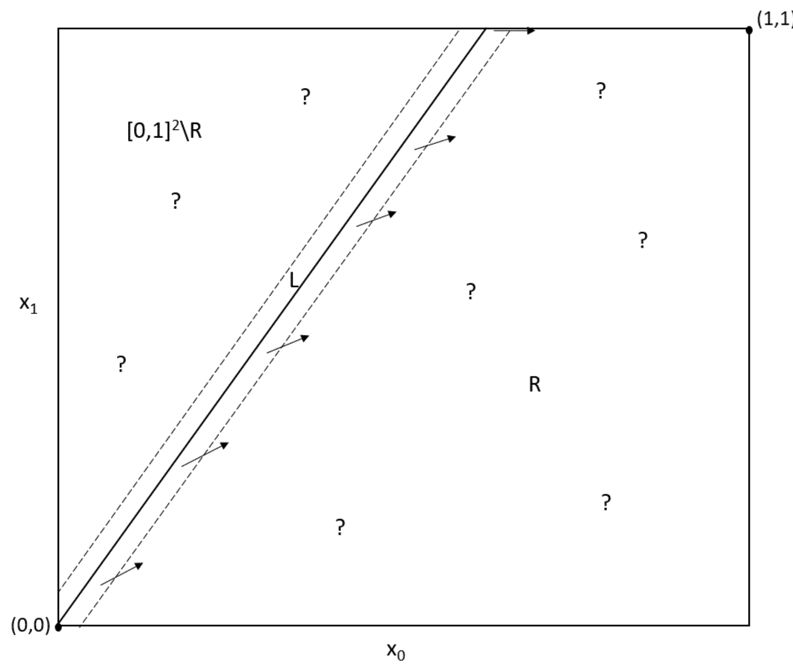

Figure 1: Phase diagram for the equilibrium population dynamic. When (x_0, x_1) is close to the line L, the population dynamic G_η moves toward the regular point $(1, 1)$. The population dynamic can therefore never exit the set R, regardless of its behavior at points away from L. Note that the figure is drawn for the cost-saving case ($\chi > \pi_1$), as the slope of L is greater than 1.

The same argument also explains why a (continuous-time) equilibrium path with a regular initial point visits only regular points. Note that regularity of the initial point is a mild requirement. If the initial generation makes its decisions on the basis of the prior alone, the initial condition will be the regular point $(x_0(0), x_1(0)) = (1, 1)$. If the initial generation also receives some exogenous informative signals about the state before making its decisions, this always results in an initial condition satisfying $x_0(0) \leq x_1(0)$, which implies regularity

[13] This argument does not depend on the fact that the point toward which the population dynamic drifts when the success rate is the same in both states is the extreme point $(1, 1)$: all that matters is that this point is regular. Theorem 1 therefore generalizes to a range of environments without this feature. See Section 6.

outside of the cost-saving innovation case. In Section 6, I show that, with exogenous signals about the state, the initial condition will also be regular even in the cost-saving case, as long as the exogenous signals are not too informative. Theorem 1 thus justifies focusing on regular states.

Theorem 1 *A regular equilibrium exists. In addition, if* (x_0, x_1) *is an equilibrium path and* $(x_0(0), x_1(0))$ *is regular, then* $(x_0(t), x_1(t))$ *is regular for all* $t \in \mathbb{R}_+$*.*

Proof. Suppose $\chi < 1$. (The $\chi = 1$ case will be established below by taking a limit.) For all $\eta \in (0,1)$, define the correspondence $G_\eta : R \rightrightarrows R$ by

$$G_\eta(x_0, x_1) = (1-\eta)(x_0, x_1) + \eta F(x_0, x_1).$$

It is clear that R is non-empty, compact, and convex. In addition, G_η is convex-valued and upper hemi-continuous for all $\eta \in (0,1)$, because F is convex-valued and upper hemi-continuous. The existence proof is thus completed by showing that G_η is non-empty (i.e., contains a point in R) for sufficiently small $\eta > 0$.

The proof that G_η is non-empty is divided into two case. First, suppose that $d((x_0, x_1), L) \geq \sqrt{2}\eta$, where L denotes the line with equation $x_1(\pi_1 - \chi) - x_0(\pi_0 - \chi) = 0$ and d denotes Euclidean distance. Then $G_\eta(x_0, x_1) \subseteq R$, so the fact that $F(x_0, x_1)$ is non-empty as a subset of $[0,1]^2$ implies that G_η is non-empty as a subset of R.

Next, I claim that, when η is sufficiently small, $F(x_0, x_1) = \{(1,1)\}$ for all (x_0, x_1) satisfying $d((x_0, x_1), L) < \sqrt{2}\eta$. This completes the proof, as $F(x_0, x_1) = \{(1,1)\}$ implies that $G_\eta(x_0, x_1) = \{(1-\eta)(x_0, x_1) + \eta(1,1)\}$, which lies in R whenever $(x_0, x_1) \in R$ (as $(1,1) \in R$ and R is convex). To show that $F(x_0, x_1) = \{(1,1)\}$, it suffices to show that $p(0; x_0, x_1) > p^*$, as this implies that $p(k; x_0, x_1) > p^*$ for all $k > 0$. By Bayes' rule, $p(0; x_0, x_1) > p^*$ if and only if

$$\left(\frac{1 - \chi - x_0(\pi_0 - \chi)}{1 - \chi - x_1(\pi_1 - \chi)}\right)^K < \frac{p}{1-p}\frac{1-p^*}{p^*}. \tag{10}$$

As $\chi < 1$, the left-hand side of (10) is uniformly continuous in (x_0, x_1) and equals 1 when $d((x_0, x_1), L) = 0$. Since $\frac{p}{1-p}\frac{1-p^*}{p^*} > 1$ (by (2)), it follows that there exists $\eta > 0$ such that

(10) holds whenever $d\left(\left(x_0, x_1\right), L\right) < \sqrt{2}\eta$.

To see that a regular equilibrium also exists when $\chi = 1$, take a sequence $\chi^n \to 1$ and a corresponding sequence of regular equilibria $(x_0^n, x_1^n, k^{*,n}, s^n) \to (x_0, x_1, k^*, s)$. Then (x_0, x_1, k^*, s) is a regular equilibrium when $\chi = 1$.

The proof that an equilibrium path with a regular initial point visits only regular points is similar and is deferred to Appendix B. ■

3 Equilibrium Efficiency: Main Results

This section contains my main results on equilibrium efficiency. Section 3.1 presents the basic result that regular equilibria always entail substantial inefficiency when the innovation is of the cost-saving type. Section 3.2 then shows that equilibrium adoption is dramatically more efficient for outcome-improving innovations than for cost-saving innovations when players' samples are large ($K \to \infty$).

The analysis makes use of some auxiliary results, which to smooth the exposition are deferred to Appendix A.

3.1 Inefficiency of Regular Equilibria with a Cost-Saving Innovation

In the cost-saving innovation case, expected welfare is bounded away from the first-best level at every regular point, and a fortiori at every regular equilibrium. In light of Theorem 1, this implies that, if the initial point is regular, then the equilibrium population dynamic exhibits substantial inefficiency at every point in time, regardless of the size of players' samples. As discussed in the introduction, the intuition for this inefficiency result is simply that with a cost-saving innovation it is efficient for failure to be *more* likely in the good state: at the first-best, players use the low-cost, high-failure innovation in the good state and use the high-cost, low-failure status quo in the bad state. But, by definition, at a regular point failure is less likely in the good state. So no regular point can be close to efficient.

In what follows, note that *expected welfare* at the point (x_0, x_1) equals

$$p\left[x_1\left(\pi_1 - c\right) + \left(1 - x_1\right)\chi\right] + \left(1 - p\right)\left[x_0\left(\pi_0 - c\right) + \left(1 - x_0\right)\chi\right]. \tag{11}$$

In particular, *first-best* expected welfare equals $p\left(\pi_1 - c\right) + \left(1 - p\right)\chi$.

Theorem 2 *Assume that $\chi > \pi_1$. Then expected welfare at every regular point is uniformly bounded away from efficiency.*

More specifically, expected welfare at every regular point falls short of the first-best by at least

$$\left(1 - p\right)\left(\frac{\chi - \pi_1}{\chi - \pi_0}\right)\left(\chi - \pi_0 + c\right) > 0. \tag{12}$$

Proof. First-best expected welfare exceeds expected welfare at the pair (x_0, x_1) by the amount

$$p\left(1 - x_1\right)\left(\pi_1 - c - \chi\right) + \left(1 - p\right)x_0\left(\chi - \pi_0 + c\right).$$

If $\chi > \pi_1$, then for any regular pair $x_0 \geq \frac{\chi - \pi_1}{\chi - \pi_0}x_1$. Hence, this excess is at least

$$p\left(1 - x_1\right)\left(\pi_1 - c - \chi\right) + \left(1 - p\right)x_1\left(\frac{\chi - \pi_1}{\chi - \pi_0}\right)\left(\chi - \pi_0 + c\right).$$

This expression is decreasing in x_1, as $p\left(\chi - \pi_0\right)\left(\pi_1 - c - \chi\right) > \left(1 - p\right)\left(\chi - \pi_1\right)\left(\chi - \pi_0 + c\right)$. (To see this, rearrange this inequality to $p\left(\pi_1 - \pi_0\right)\left(-c\right) > \left(\chi - \pi_1\right)\left(\chi - \pi_0 + c\right)$, and note that $-c > \chi - \pi_1$ and $p\left(\pi_1 - \pi_0\right) > \chi - \pi_0 + c$, by (1).) Therefore, inefficiency is minimized at $x_1 = 1$, which gives (12). ■

To interpret (12), note that the proof shows that the minimal inefficiency at a regular point is achieved at the point $(x_0, x_1) = \left(\frac{\chi - \pi_1}{\chi - \pi_0}, 1\right)$.[14] At this point, the correct action is always taken in state 1, and the wrong action is taken in state 0 with probability $\frac{\chi - \pi_1}{\chi - \pi_0}$. As the loss from taking the wrong action in state 0 is $\chi - \pi_0 + c$, it follows that the minimal inefficiency at any regular point is $\left(1 - p\right)\left(\frac{\chi - \pi_1}{\chi - \pi_0}\right)\left(\chi - \pi_0 + c\right)$.

To get a sense of the magnitude of this loss, note that a completely uninformed player who follows the prior and always takes action 1 suffers a loss of $\left(1 - p\right)\left(\chi - \pi_0 + c\right)$. Thus,

[14] Somewhat surprisingly, this holds regardless of the prior p.

the ratio of the minimal loss at a regular point to the loss of an uninformed player equals $\frac{\chi-\pi_1}{\chi-\pi_0}$. For instance, if $\chi = 1$ and $\pi_0 = 0$, this minimum loss ratio ratio equals $1-\pi_1$: that is, at most fraction π_1 of the loss suffered by a completely uninformed player can be eliminated by social learning. More generally, note that the loss ratio is increasing in χ and π_0 but decreasing in π_1. In this sense, social learning is potentially more powerful when the status quo is less effective and when outcomes under the innovation are less noisy.

Finally, in the pure cost-saving innovation case ($\chi = 1$), the minimal inefficiency at a regular point (12) goes to 0 as $\pi_1 \to 1$. However, results in Section 4 demonstrate that, when K is not too large, substantial inefficiency persists even as $\pi_1 \to 1$ at any regular *equilibrium* point. Thus, (12) is not a tight bound on equilibrium inefficiency in the pure cost-saving case. (In contrast, Proposition 6 below shows that equilibrium in the pure outcome-improving case tends toward efficiency as $\pi_0 \to 0$ and $\pi_1 \to 1$.)

3.2 Many Observations: Learning vs. Confounding

The inefficient adoption of cost-saving innovations documented by Theorems 1 and 2 applies regardless of the size of players' samples. Conversely, outcome-improving innovations are adopted efficiently as $K \to \infty$. There is thus a dramatic difference in efficiency between the equilibrium adoption of cost-saving innovations and outcome-improving innovations when players observe large samples.

To see why outcome-improving innovations are adopted efficiently, note that Bayesian rationality implies that the adoption rate in the good state is greater than that in the bad state and is also bounded away from 0: formally, Proposition 13 in Appendix A establishes that $x_1 \geq \max\left\{x_0, \frac{p-p^*}{1-p^*}\right\}$ in any equilibrium. This in turn implies that the success rate cannot be too similar in the two states, as if $\chi < \pi_1$ then

$$\begin{aligned} x_1(\pi_1-\chi) - x_0(\pi_0-\chi) &\geq x_1(\pi_1-\chi) - x_0 \max\{\pi_0-\chi, 0\} \\ &\geq x_1 \min\{\pi_1-\pi_0, \pi_1-\chi\} \\ &\geq \frac{p-p^*}{1-p^*} \min\{\pi_1-\pi_0, \pi_1-\chi\} > 0. \end{aligned} \tag{13}$$

Given (13), players learn the state as $K \to \infty$ by the law of large numbers, and efficiency follows.

Theorem 3 *Assume that $\chi < \pi_1$. Fix any sequence of equilibria $(x_{0,K}, x_{1,K}, k^*_K, s_K)$ indexed by K. Then $\lim_{K\to\infty} x_{0,K} = 0$ and $\lim_{K\to\infty} x_{1,K} = 1$.*

Proof. See Appendix B. ■

On the other hand, Theorems 1 and 2 show that there exists an upper bound on the share of players taking the correct action in the cost-saving case, independently of K. How is it possible that players fail to learn the state as $K \to \infty$? The only explanation is that each individual observation becomes completely uninformative in the limit: that is, the success rate becomes the same in both states.[15] This is a striking example of confounded learning.[16]

Proposition 3 *Assume that $\chi > \pi_1$. Fix a sequence of regular equilibria $(x_{0,K}, x_{1,K}, k^*_K, s_K)$ indexed by K. Then*

$$\lim_{K\to\infty} \frac{x_{0,K}(\pi_0 - \chi)}{x_{1,K}(\pi_1 - \chi)} = 1.$$

Proof. Suppose toward a contradiction that there exists $\varepsilon > 0$ and a sequence $(K_n)_{n=1}^{\infty}$ such that $x_{0,K_n}(\pi_0 - \chi) < (1+\varepsilon)\, x_{1,K_n}(\pi_1 - \chi)$ for all n. Then, for all n,

$$x_{1,K_n}(\pi_1 - \chi) - x_{0,K_n}(\pi_0 - \chi) > \varepsilon x_{1,K_n}(\chi - \pi_1) \geq \varepsilon \frac{p - p^*}{1 - p^*}(\chi - \pi_1) > 0.$$

This gives a uniform lower bound on $x_{1,K_n}(\pi_1 - \chi) - x_{0,K_n}(\pi_0 - \chi)$, so the law of large numbers (more precisely, the same argument as in the proof of Theorem 3) implies that $\lim_{n\to\infty} x_{0,K_n} = 0$ and $\lim_{n\to\infty} x_{1,K_n} = 1$. As $\chi > \pi_1$, this contradicts the assumption that (x_{0,K_n}, x_{1,K_n}) is regular for all n. ■

Thus, with an outcome-improving innovation, players learn the state and adopt the innovation efficiently whenever they can observe the outcomes of sufficiently many of their peers. In contrast, with a cost-saving innovation, the possibility of confounded learning prevents efficient adoption of the innovation, even though learning is never completely confounded in equilibrium for any finite sample size.

[15] In other words, equilibria approach the line L in Figure 1.

[16] Other well-known examples include McLennan (1984), Kiefer (1989), Piketty (1995), and Smith and Sørensen (2000). A difference is that confounded learning occurs in equilibrium in these papers, while in my model it arises only asymptotically as $K \to \infty$.

Theorems 1, 2, and 3 yield the main conclusion of this paper: as $K \to \infty$, outcome-improving innovations are adopted efficiently, but cost-saving innovations are not. The key economic intuition is that it is hard to learn about cost-saving innovations, because it is not clear if observing good outcomes is good news or bad news.

4 Equilibrium Efficiency: Small Samples

Are outcome-improving innovations also adopted more efficiently than cost-saving innovations when K is small? As outcomes are noisy, there is no hope that equilibria with fixed K can be efficient unless the outcome under the innovation is close to deterministic ($\pi_0 \to 0$ and/or $\pi_1 \to 1$), and thus either the pure outcome-improving or pure cost-saving case applies. I therefore focus on the question of whether equilibria tend toward efficiency in these limits in the $K = 1$ case (Section 4.1) and the $K > 1$ case (Section 4.2). I again find clear distinctions between the outcome-improving and cost-saving cases. For example, if $K > 1$ then equilibria tend toward efficiency as $\pi_0 \to 0$ and $\pi_1 \to 1$ in the pure outcome-improving case but remain persistently inefficient in the pure cost-saving case.

4.1 One Observation: Under- vs. Over-Adoption

When $K = 1$, the equilibrium does not tend toward efficiency as the risky action becomes close to deterministic (i.e., π_0 approaches 0 and π_1 approaches 1). More interestingly, the nature of the inefficiency is different in the pure outcome-improving innovation case ($\chi = 0$) and in all other cases ($\chi > 0$). In the pure outcome-improving innovation case, in the unique equilibrium the innovation is fully rejected in the bad state, but it is not fully adopted in the good state. The innovation is thus *under-adopted*. On the other hand, outside of this case, the innovation is fully adopted in the good state but is not fully rejected in the bad state: that is, it is *over-adopted*.

Proposition 4 *Assume each player observes only one other player's outcome ($K = 1$). There is a unique equilibrium. Moreover,*

1. *In the pure outcome-improving innovation case, for all $\varepsilon > 0$ there exists $\delta > 0$ such*

that, if $1 > \pi_1 > 1 - \delta$*, then in equilibrium* $x_0 < \varepsilon$ *and* $x_1 < \frac{p-\hat{p}}{p(1-\hat{p})} + \varepsilon$*, where* $\hat{p} := \frac{\chi-\pi_0+c}{1-\pi_0}$*. The innovation is thus under-adopted.*

2. *Outside of the pure outcome-improving innovation case, for all* $\varepsilon > 0$ *there exists* $\delta > 0$ *such that, if* $\pi_1 > 1-\delta$*, then in equilibrium* $x_0 = \frac{\chi}{1-\pi_0+\chi}$ *and* $x_1 > 1-\varepsilon$*. The innovation is thus over-adopted.*

Proof. The proofs of this and all subsequent results in the paper are deferred to Appendix B. ■

To illustrate Proposition 4 in the context of an example, consider an agricultural community that relies on planting traditional maize, and compare the resulting adoption patterns when the community faces the introduction of a high-yield variety of maize and when it faces the introduction of a completely new crop such as pineapple.[17] In the high-yield maize case, the bad outcome corresponds to a small maize harvest and the good outcome corresponds to a large maize harvest; in the pineapple case, the bad outcome is maize but no pineapple, while the good outcome is maize and pineapple. Suppose that in both cases the new crop is almost certainly successful if the underlying soil conditions in the community are favorable.

The critical distinction between the cases where the new crop is high-yield maize and where it is pineapple is that the status quo crop (traditional maize) might produce the good outcome (a large maize harvest) in the former case, but can never produce the good outcome (maize and pineapple) in the latter. Proposition 4 then says that, if each farmer observes only one neighbor's harvest before planting her own crop, the pattern of technology adoption will be completely different in the two cases. In the high-yield maize case, all farmers will (profitably) adopt the new crop when the soil conditions are favorable, but some farmers will (unprofitably) do the same even when the soil is unfavorable. In the pineapple case, no farmers will mistakenly adopt the pineapple when the soil is unfavorable, but some farmers will fail to adopt when the soil is favorable. High-yield maize is thus over-adopted, while pineapple is under-adopted.

[17] The example of maize and pineapple is inspired by the influential study of Conley and Udry (2010), though not all aspects of their environment fit my model. Most importantly, they assume that farmers observe their neighbors' planting decisions and fertilizer utilization in addition to their crop yields.

What is the intuition for this result? In the high-yield maize case, a farmer will occasionally observe a large maize harvest even if the soil is bad. Such an observation constitutes good news about the soil (as when $K = 1$ the steady state is always regular). Farmers who see large maize harvests will therefore plant high-yield maize, and this behavior leads to a positive rate of high-yield maize adoption even when the soil is bad.

When instead the soil is good, the high-yield variety almost always produces a large harvest (as π_1 is assumed to be close to 1), and in particular it produces a small harvest with a much smaller probability than that which with the traditional variety produces a large harvest. Therefore, under the equilibrium strategy of planting the high-yield variety after seeing a large harvest and planting the traditional variety after seeing a small harvest, in a steady state almost all the farmers must plant the high-yield variety.

Next, consider the pineapple case. The key observation is that farmers must now mix between the two varieties after observing "no-pineapple." For, if farmers never planted pineapple after observing no-pineapple, then the fact that maize seed cannot produce pineapple would imply that the pineapple planting rate would gradually decline to 0 even when the soil is good (whenever $\pi_1 < 1$ and $K = 1$), which cannot happen. Now, for the farmers to be willing to mix after observing no-pineapple, the pineapple planting rate must be set so that this observation constitutes just bad enough news about the soil condition to make them indifferent. This requirement implies that the planting rate is bounded away from 1 when the soil is good. Finally, as $\pi_1 \to 1$ the planting rate can remain bounded away from 1 when the soil is good only if the probability of planting pineapple after observing no-pineapple goes to 0. This in turn implies that the planting rate converges to 0 when the soil is bad, as is efficient.

More generally, the key prediction of Proposition 4 is that qualitatively new technologies—those that can produce results that cannot be confused with results coming from existing technologies—are under-adopted, while new technologies that are only more likely to produce good results are over-adopted.

Another interesting consequence of Proposition 4 is that, if the ex ante expected value of a pure outcome-improving innovation is only slightly positive (i.e., p is only slightly greater than p^*), then in equilibrium very few players adopt the innovation even when it is good.

(This follows because $x_1 \to \frac{p-\hat{p}}{p(1-\hat{p})}$ and $|\hat{p} - p^*| \to 0$ as $\pi_1 \to 1$.) The intuition is that equilibrium requires players to be indifferent to the innovation after observing a failure, and when the ex ante value of the innovation is small this requires that observing a failure is a very weak signal about the quality of the innovation. But observing a failure can be a very weak signal only if very few people are using the innovation. Hence, pure outcome-improving innovations that yield small expected benefits over the status quo can never be widely adopted, even if the realized benefits are in fact large.

I conclude the analysis of the $K = 1$ case by discussing the effects of the model parameters (χ, π_0, π_1, p, and c) on equilibrium welfare.

Proposition 5 *When $K = 1$, expected welfare is strictly increasing in π_0, π_1, and p, and is strictly decreasing in c. Expected welfare is non-decreasing in χ, but is constant in χ in the region where the equilibrium is in mixed strategies.*[18]

The most interesting part of the result is that expected welfare does not increase with χ when the equilibrium is in mixed strategies. In this region, players take the safe action with positive probability in both states, and yet improving the payoff to the safe action has no effect on their expected payoffs. The explanation is that, in a mixed equilibrium, players experiment after observing success and mix after observing failure, so expected welfare must be the same as it would be if players always experimented, and this payoff is independent of the payoff to the safe action. A further observation is that the safe action is played with different probabilities in the two states, so changing χ does affect players' payoffs conditional on the state. In particular, one can check that improving the payoff to the safe action actually *decreases* welfare when the state is good. The mechanism behind this counterintuitive comparative static is that, when χ increases, observing failure must be less informative about the state for players to be willing to experiment after observing failure. But this implies that x_1 must be lower, which reduces welfare in state 1 (and in particular more than offsets the direct benefit of increasing χ).

[18] The equilibrium is in mixed strategies if and only if inequality (17) in Appendix B fails to hold.

4.2 Multiple Observations: Efficiency vs. Persistent Inefficiency

The inefficiency in the $K = 1$ case documented in Proposition 4 vanishes for all $K > 1$ in the pure outcome-improving case, while in the pure cost-saving case inefficiency diminishes only slowly as K increases. Thus, as in the $K \to \infty$ case, equilibrium adoption is dramatically more efficient with an outcome-improving innovation.

Proposition 6 *Assume that $K > 1$.*

1. *In the pure outcome-improving innovation case, for all $\varepsilon > 0$ there exists $\delta > 0$ such that, if $\pi_0 < \delta$ and $\pi_1 > 1-\delta$, then in any equilibrium $x_0 < \varepsilon$ and $x_1 > 1-\varepsilon$. Adoption is thus approximately efficient.*

2. *In the pure cost-saving innovation case, for all $\varepsilon > 0$ there exists $\delta > 0$ such that, if $\pi_0 < \delta$ and $\pi_1 > 1 - \delta$, then in any regular equilibrium $\left| x_0 - (1 - x_0)^K \right| < \varepsilon$ and $x_1 > 1 - \varepsilon$. Over-adoption thus persists for every finite K.*

   In contrast, for all $\varepsilon > 0$ there exist $\delta > 0$ and $\bar{K} > 0$ such that, if $\pi_0 < \delta$, $\pi_1 > 1 - \delta$, and $K > \bar{K}$, then in any regular equilibrium $x_0 < \varepsilon$ and $x_1 > 1 - \varepsilon$. Adoption is thus approximately efficient when K is large.[19]

In the pure outcome-improving case, the stark difference between the $K = 1$ and $K > 1$ cases may be explained by considering the equation for x_1 that results when $x_0 = 0$ and players experiment if and only if they observe at least one success:

$$x_1 = 1 - (1 - x_1\pi_1)^K .$$

When $K = 1$, for any $\pi_1 < 1$ the unique solution to this equation is $x_1 = 0$. In contrast, for any $K > 1$, when π_1 is sufficiently large this equation also admits a positive solution, and this solution converges to 1 as $\pi_1 \to 1$. Thus, when $K = 1$, equilibrium requires players to

[19] There is of course no contradiction with Theorem 2: in the cost-saving case, for any values of π_0 and π_1 efficiency in any regular equilibrium is uniformly bounded away from efficiency for all K, but the degree of inefficiency vanishes as $\pi_0 \to 0$ and $\pi_1 \to 1$. Thus, in the pure outcome-improving case, equilibrium is efficient in the $K \to \infty$ limit *as well as* in the $\pi_0 \to 0/\pi_1 \to 1$ limit, while in the pure cost-saving case equilibrium is efficient only if *both* K is large and π_0 is small/π_1 is large.

mix after observing failure (which then imposes an upper bound on x_1), while when $K > 1$ it is possible for x_1 to be close to 1 even in a pure strategy equilibrium.

In the pure cost-saving case, when $\pi_0 = 0$ the population share using the innovation in the bad state satisfies the equation $x_0 = (1 - x_0)^K$. The solution to this equation converges to 0 rather slowly as $K \to \infty$. For example, if $K = 5$ then $x \approx 0.25$, and if $K = 50$ then $x \approx 0.06$. Inefficiency is therefore substantially greater in the pure cost-saving case than in the pure outcome-improving case, even when K is relatively large.

A final observation concerns the comparative statics of expected welfare with respect to K. While Proposition 6 illustrates settings where larger samples lead to higher equilibrium welfare (as one would expect), this comparative static does not hold in general. Indeed, Example 2 in Appendix A shows that welfare does not always unambiguously increase when players observe larger samples even if one restricts attention to stable regular equilibria.

5 Convergence to Equilibrium in the $K = 1$ Case

A question left unanswered by the analysis so far is whether the equilibrium population dynamic always converges to a steady state. The existence of multiple stable regular steady states—as illustrated in Example 2 in Appendix A—suggests that this may be a difficult question to resolve, and I have not been able to provide a complete answer. However, convergence to the unique steady state does always occur in the $K = 1$ case. Interestingly, this convergence may not be monotonic.

The convergence result for the $K = 1$ case is as follows.

Proposition 7 *When $K = 1$, the unique equilibrium (x_0^*, x_1^*) is globally asymptotically stable: for any equilibrium path (x_0, x_1), $\lim_{t\to\infty} (x_0(t), x_1(t)) = (x_0^*, x_1^*)$.*

A simple example illustrates that convergence can be non-monotone.

Proposition 8 *When $K = 1$, $\chi = 0$, and $\pi_0 = 0$, there is a unique equilibrium path starting from the initial condition $(x_0(0), x_1(0)) = (1, 1)$, and the population dynamics are as follows:*

1. *$x_1(t)$ decreases at rate $1 - \pi_1$ until reaching its steady-state value x_1^* at some finite time T, and then remains constant forever.*

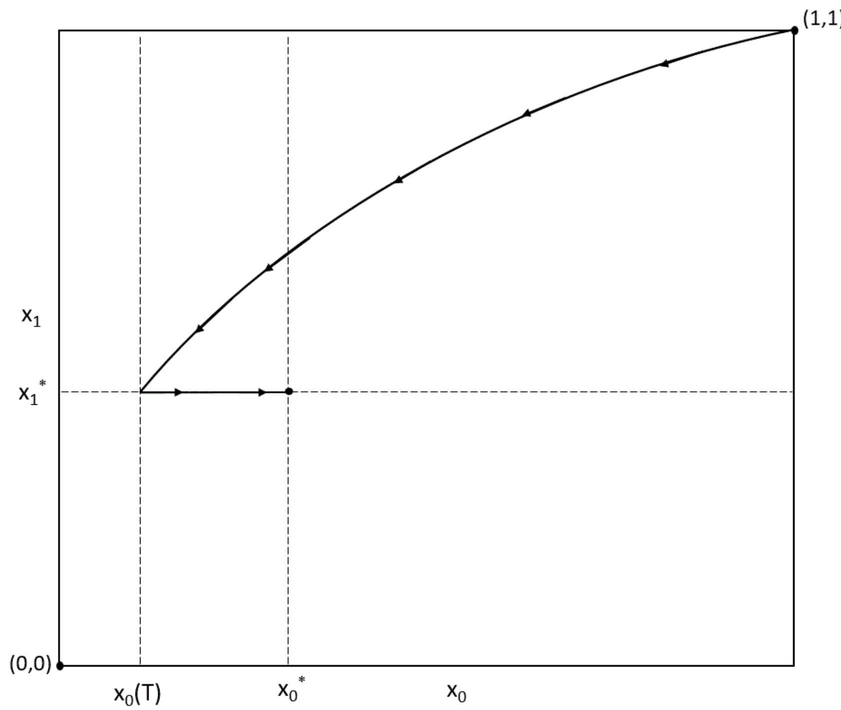

Figure 2: Population dynamics with $K = 1$, $\chi = 0$, $\pi_0 = 0$, initial condition $(x_0(0), x_1(0)) = (1, 1)$, and $x_0(T) < x_0^*$.

2. *$x_0(t)$ decreases at rate 1 up to time T, and then converges monotonically to its equilibrium value x_0^*. The overall path $x_0(t)$ is thus non-monotone if and only if $x_0(T) < x_0^*$. Finally, there exists $\delta > 0$ such that if $\pi_1 > 1 - \delta$ then $x_0(T) < x_0^*$.*

Figure 2 illustrates the equilibrium path in the case where $x_0(T) < x_0^*$. The intuition for why the path is non-monotone is as follows: First, $x_1(t)$ must decrease over time until reaching its steady-state value x_1^*, as if $x_1(t) > x_1^*$ then observing failure is so informative that a player who observes failure will never use the innovation. Once $x_1(t)$ reaches x_1^*, observing failure becomes just informative enough that a player who observes failure is indifferent between the innovation and the status quo. At this point, players who observe failure mix, using the innovation with the probability that keeps the fraction experimenting in the good state constant. Finally, if the fraction experimenting in the bad state is sufficiently low by the time $x_1(t)$ reaches x_1^* (i.e., if $x_0(T) < x_0^*$), then this probability of experimentation is high enough that the share experimenting in the bad state starts to increase, before eventually approaching its steady state level.

In addition to featuring non-monotone population dynamics, this example also has the interesting property that the well-known *improvement principle* of Banerjee and Fudenberg does not hold. The improvement principle states that average welfare in the population is non-decreasing over time. The improvement principle always holds when players observe a random sample of their predecessors' *actions* (as in Banerjee and Fudenberg), as each member of the newest generation can guarantee herself the average level of welfare simply by copying one of the actions she observes at random. This argument does not apply when players only observe their predecessor's outcomes. Indeed, it is clear that the improvement principle fails in the setting of Proposition 8 whenever population dynamics are non-monotone: this follows because, after time T, average welfare in state 1 is constant while average welfare in state 0 is strictly decreasing.

While I have not resolved the question of whether the continuous-time equilibrium population dynamic always converges to an equilibrium—or whether instead infinite cycles are possible—the corresponding discrete-time population dynamic can involve cycles. For example, consider the discrete-time game where in each period $t = 0, 1, 2, \ldots$, a continuum of players of mass 1 take actions, and before acting each player observes a random sample of K outcomes from the previous period.[20] Assuming that all players in the first generation follow the prior and take action 1, in the simple case where $K = 1$ and $\pi_0 = 0$ the equilibrium is generically unique and corresponds to the unique equilibrium in Acemoglu and Wolitzky (2014): the fraction of players taking action 1 in period t in the good state equals $\pi_1^{t \bmod T}$, where

$$T = \left\lceil \frac{\log\left(1 - \frac{1-p}{p}\frac{p^*}{1-p^*}\right)}{\log \pi_1} \right\rceil > 1$$

is the first time at which a player's posterior after observing a failure lies above p^*.[21] The equilibrium population dynamic thus follows a deterministic cycle with period T.

[20]In the notation of the proof of Theorem 1, this model corresponds to the case $\eta = 1$, where the entire population turns over every period.

[21]The fraction of players taking action 1 in period t in the bad state equals 1 if $t = 0 \bmod T$ and 0 otherwise.

6 Extensions

This section extends the model in two directions. Section 6.1 considers more general physical environments, which allow for a continuum of actions, heterogeneous players, and multiple states of the world. Section 6.2 considers more general information structures, which allow for additional signals of the state (or other variables measurable with respect to the state, such as the aggregate adoption rate) and additional signals of others' actions. The emphasis in all cases is on assessing the robustness of the main results, Theorems 1, 2, and 3.

6.1 More General Physical Environments

6.1.1 Continuous Actions and Convex Costs

In the model so far, each player faces a binary choice of whether or not to adopt the innovation. In some settings, it is more realistic to model the intensity of adoption as a continuous choice, where the cost of adoption is convex in intensity. For example, if the innovation corresponds to a high-yield crop variety that is also more labor-intensive than the traditional variety, then a standard labor-leisure tradeoff would yield an adoption cost convex in the share of one's land planted with the high-yield variety. Alternatively, if high-yield seeds are more expensive than traditional seeds and players are credit constrained, then adoption costs would again be convex in the intensity of adoption. Note also that the assumption that actions are unobservable may be a more reasonable approximation in a continuous-action model: it may be easy for a farmer to tell if her neighbor switches his entire crop from the traditional to the high-yield variety, but hard for her to tell exactly what fraction of his land is planted with each variety.

Such a model would be nearly isomorphic to the one presented above. To see this, suppose that $X = [0, 1]$ and a player's payoff from taking action x in state θ is

$$x\pi_\theta + (1 - x)\,\chi - c\,(x)\,,$$

where $c : \mathbb{R} \to \mathbb{R}$ is a continuously differentiable and strictly convex function. Assume that, before taking her action, a player who enters the population at time t observes the outcomes of K independent Bernoulli trials with parameter

$$x_\theta(t)\,\pi_\theta + (1 - x_\theta(t))\,\chi,$$

where $x_\theta(t)$ the average action taken among players in the population at time t. The interpretation is now that each player chooses whether to use the innovation or the status quo for each of a continuum of "instances," and that before making her choice she observes the outcome of a random sample of K instances. For example, a farmer may have to decide whether to use a new crop variety on each of many plots of land, or a doctor may decide whether to use a new drug on each of many patients; and players may observe the outcomes of a random sample of plots or patients.[22]

Let

$$x^*(\tilde{p}) = (c')^{-1}(\tilde{p}\pi_1 + (1-\tilde{p})\,\pi_0 - \chi)$$

be the optimal action of a player for whom $\Pr(\theta = 1) = \tilde{p}$. The analogues of assumptions (1) and (2) above—which imply that always playing $x = 0$ or $x = 1$ is not an equilibrium—are $x^*(p) > 0$ and $x^*(p(0;1,1)) < 1$, where $p(k;x_0,x_1)$ is defined in equation (3). The resulting population trajectory is given by

$$\dot{x}_\theta(t) = \sum_{k=0}^{K} \phi_\theta(k; x_\theta(t))\, x^*(p(k; x_0(t), x_1(t))) - x_\theta(t) \text{ for } \theta = 0, 1.$$

Theorems 1, 2, and 3 easily extend to this continuous-action environment. In particular, if the average actions $(x_0(t), x_1(t))$ are close to the line L with equation $x_1(\pi_1 - \chi) - x_0(\pi_0 - \chi) = 0$, then $x^*(p(k; x_0(t), x_1(t)))$ is close to $x^*(p)$ for all k, and therefore the population dynamic drifts toward the regular point $(x^*(p), x^*(p))$. This implies that the

[22] Note that the assumption that a player observes K independent instances rules out situations where, say, a player randomly samples one other player and observes K of her instances. The independence assumption here is thus rather restrictive. However, without this assumption one would have to keep track of the entire distribution of actions in each state rather than only the average action, which would likely make the model intractable. A possible interpretation of the independence assumption is that the outcomes of all instances are recorded in an anonymized database and each player observes a random sample of outcomes from the database. An alternative interpretation is that only one instance for each player is "visible"—for example, the farmer's plot of land that is located closest to a main road—and each player observes the visible instance of a random sample of other players.

population dynamic can never exit the regular region R.

6.1.2 Heterogeneous Players

Another stylized assumption in the baseline model is that all players are identical. It would be more realistic to assume that players differ in exogenous characteristics that affect their inclinations to adopt the innovation. For example, farmers have information about features of their own soil, and it may be known that, if a new crop variety is effective, it will be particularly effective for certain types of soil. Whether enriching the model on this dimension affects the main results turns out to hinge on whether players' characteristics are observable—for example, if a farmer knows the soil characteristics of the neighbors whose harvests she observes. In particular, unobservable heterogeneity does not affect the results, while observable heterogeneity will generically lead to efficient adoption even in the cost-saving case.

Formally, suppose there are $|Q|$ types of players. Each player knows her own type, and the (constant) population share of players of type q is σ_q. A player's type affects her probability of success when using the innovation, in that a player of type q who uses the innovation in state θ succeeds with probability $\pi_{\theta,q}$. Assume that assumptions (1) and (2) are satisfied for each type q.

Suppose first that each player continues to observe a random sample of K other players' outcomes, without observing their types. The population dynamic must now keep track of the share of players of each type taking each action. Thus, an equilibrium path consists of a list of functions

$$\left(x_{0,q} : \mathbb{R}_+ \to [0,1]\,, x_{1,q} : \mathbb{R}_+ \to [0,1]\,, k_q^* : \mathbb{R}_+ \to \{0, \dots, K\}\,, s_q : \mathbb{R}_+ \to [0,1]\right)_{q \in Q},$$

a vector of population shares $\left(x_{0,q}, x_{1,q}\right)_{q \in Q}$ is regular if and only if

$$\sum_{q \in Q} \sigma_q x_{1,q} \left(\pi_{1,q} - \chi\right) \geq \sum_{q \in Q} \sigma_q x_{0,q} \left(\pi_{0,q} - \chi\right), \tag{14}$$

the probability that a new player observes k successes when the state is θ and fraction x_q of the type-q players in the population use the innovation is

$$\hat{\phi}_\theta\left(k;(x_q)_{q\in Q}\right)=\begin{pmatrix}K\\k\end{pmatrix}\left[\chi+\sum_{q\in Q}\sigma_q x_{\theta,q}\left(\pi_{\theta,q}-\chi\right)\right]^k\left[1-\chi-\sum_{q\in Q}\sigma_q x_{\theta,q}\left(\pi_{\theta,q}-\chi\right)\right]^{K-k},$$

the population trajectory is given by

$$\dot{x}_{\theta,q}(t)=\left\{\begin{array}{c}\hat{\phi}_\theta\left(k_q^*(t);(x_{\theta,q}(t))_{q\in Q}\right)s_q(t)+\sum_{k=k_q^*(t)+1}^K\hat{\phi}_\theta\left(k;(x_{\theta,q}(t))_{q\in Q}\right)-x_{\theta,q}(t)\\ \text{if (14) holds at }\left(x_{0,q}(t),x_{1,q}(t)\right)_{q\in Q};\\ \hat{\phi}_\theta\left(k_q^*(t);(x_{\theta,q}(t))_{q\in Q}\right)s_q(t)+\sum_{k=0}^{k_q^*(t)-1}\hat{\phi}_\theta\left(k;(x_{\theta,q}(t))_{q\in Q}\right)-x_{\theta,q}(t)\\ \text{if (14) fails at }\left(x_{0,q}(t),x_{1,q}(t)\right)_{q\in Q}\end{array}\right\},$$

and a player's assessment of the probability that $\theta=1$ after observing k successes when the population shares are given by $\left(x_{0,q},x_{1,q}\right)_{q\in Q}$ equals

$$\begin{aligned}&p\left(k;\left(x_{0,q},x_{1,q}\right)_{q\in Q}\right)\\ =&\left[1+\frac{1-p}{p}\frac{\left[\chi+\sum_{q\in Q}\sigma_q x_{0,q}\left(\pi_{0,q}-\chi\right)\right]^k\left[1-\chi-\sum_{q\in Q}\sigma_q x_{1,q}\left(\pi_{0,q}-\chi\right)\right]^{K-k}}{\left[\chi+\sum_{q\in Q}\sigma_q x_{1,q}\left(\pi_{1,q}-\chi\right)\right]^k\left[1-\chi-\sum_{q\in Q}\sigma_q x_{1,q}\left(\pi_{1,q}-\chi\right)\right]^{K-k}}\right]^{-1}.\end{aligned}$$

Although the vector of population shares now lies in $\mathbb{R}^{2|Q|}$, Theorems 1, 2, and 3 generalize immediately to this setting. For Theorem 1, the argument is simply that the boundary of the region of regular points defined by the hyperplane with equation $\sum_{q\in Q}\sigma_q x_{1,q}\left(\pi_{1,q}-\chi\right)-\sum_{q\in Q}\sigma_q x_{0,q}\left(\pi_{0,q}-\chi\right)=0$ separates this region from the efficient point $\left(x_{0,q}=0,x_{1,q}=1\right)_{q\in Q}$, and, whenever $\left(x_{0,q},x_{1,q}\right)_{q\in Q}$ is close to this hyperplane, $p\left(k;\left(x_{0,q},x_{1,q}\right)_{q\in Q}\right)$ is close to p for all $k\in\{0,\ldots,K\}$. Thus, so long as assumption (1) is satisfied for each type q, the population dynamic will drift toward to the regular point $\left(x_{0,q}=1,x_{1,q}=1\right)_{q\in Q}$ whenever it is close enough to the boundary of the regular region, and the dynamic can therefore never exit the regular region. The generalizations of Theorems 2 and 3 are also immediate.

However, matters are quite different when each player instead observes a random sample of K other players' outcomes *and types*. In this case, as long as the ratio $\left(\pi_{1,q}-\chi\right)/\left(\pi_{0,q}-\chi\right)$ is not the same for all types q, when K is sufficiently large the equilibrium population dynamic with initial point $(x_{0,q} = 1, x_{1,q} = 1)_{q \in Q}$ converges close to the efficient point $(x_{0,q} = 0, x_{1,q} = 1)_{q \in Q}$ in both the outcome-improving and cost-saving cases. The intuition is that, when $x_{1,q}(\pi_{1,q} - \chi)$ hits $x_{0,q}(\pi_{0,q} - \chi)$ for some type q, new players can still infer the state with high probability solely on the basis of the outcomes of players with types other than q. So long as $(\pi_{1,q} - \chi) / (\pi_{0,q} - \chi)$ is not the same for all q, the outcomes of players of some type will always be informative of the state at each point on the line connecting the initial point $(x_{0,q} = 1, x_{1,q} = 1)_{q \in Q}$ to the efficient point $(x_{0,q} = 0, x_{1,q} = 1)_{q \in Q}$.

6.1.3 Multiple States

The issue of whether and how the main results extend with more than two states of the world is slightly subtle. I therefore provide a more detailed analysis for this case than for the other extensions of the model. I find that, (i) unlike in the two-state case, efficient adoption of a cost-saving innovation starting from a regular point is sometimes possible with multiple states, but (ii) for a large set of parameters, this does not occur, and (iii) efficient adoption of a cost-saving innovation necessarily involves "complex" Bayesian updating—where the information content of a sample is not monotonic in the number of observed successes—and is thus arguably somewhat implausible.

To start with an example, consider the cost-saving case, and suppose there is a third, "very good" state, $\theta = 2$, with $\pi_0 < \pi_1 < \pi_2 < \chi$. Denote the prior probability of state θ by p_θ, and continue to assume that the risky action is optimal under the prior: $p_0\pi_0 + p_1\pi_1 + p_2\pi_2 - c > \chi$.

Can there be an equilibrium path leading from the initial point $(x_0, x_1, x_2) = (1, 1, 1)$ to the efficient point $(0, 1, 1)$? The answer depends on whether the risky action is optimal under the prior, conditional on the state lying in the set $\{0, 1\}$ or $\{0, 2\}$: that is, on whether $\frac{p_0}{p_0+p_1}\pi_0 + \frac{p_1}{p_0+p_1}\pi_1 - c$ and $\frac{p_0}{p_0+p_2}\pi_0 + \frac{p_2}{p_0+p_2}\pi_2 - c$ are greater or less than χ. To see why, note that any equilibrium path leading from $(1, 1, 1)$ to $(0, 1, 1)$ must pass through a point (x_0, x_1, x_2) with $x_0(\pi_0 - \chi) = x_1(\pi_1 - \chi)$ and a point (x_0', x_1', x_2') with $x_0'(\pi_0 - \chi) = x_2'(\pi_2 - \chi)$. In the first case (say), after every sample the ratio of probability weights assigned to states 0 and 1 under the posterior will be the same as the ratio under the prior. So if the risky action is optimal under the prior (conditional on the state lying in $\{0, 1\}$), then the risky action will be optimal after observing any sample at the point (x_0, x_1, x_2), which implies that the equilibrium path can never reach a point with $x_0(\pi_0 - \chi) > x_1(\pi_1 - \chi)$. On the other hand, if $\chi > \max\left\{\frac{p_0}{p_0+p_1}\pi_0 + \frac{p_1}{p_0+p_1}\pi_1, \frac{p_0}{p_0+p_2}\pi_0 + \frac{p_2}{p_0+p_2}\pi_2\right\} - c$, then an equilibrium path from $(1,1,1)$ to $(0,1,1)$ may exist for sufficiently large K.

To address this possibility, I will present two results that provide conditions guaranteeing that the main results from the two-state model generalize to the case of multiple states.

Suppose there are $n+1$ states, $\Theta = \{0, \dots, n\}$, with corresponding conditional success rates $\pi_0 < \dots < \pi_n$ and prior probabilities $(p_0, \dots, p_n)$. Assume $\pi_0 - c < \chi < \pi_n - c$ and let $\theta^* \in \{0, \dots, n-1\}$ satisfy $\pi_{\theta^*} - c < \chi < \pi_{\theta^*+1} - c$, so that θ^* is the best state at which the innovation is still inefficient. Say that an *asymptotically efficient path exists* if there exists a sequence of equilibrium paths $\left(x_0^K, \dots, x_n^K\right)$ indexed by K such that $\left(x_0^K(0), \dots, x_n^K(0)\right) = (1, \dots, 1)$ and

$$\lim_{K\to\infty}\lim_{t\to\infty}\left(x_0^K(t), \dots, x_n^K(t)\right) = \left(\underbrace{0, \dots 0}_{\theta^* \text{ times}}, \underbrace{1, \dots 1}_{n-\theta^*+1 \text{ times}}\right).$$

Theorems 1 and 2 imply that, in the two-state case $(n = 1)$, an asymptotically efficient path fails to exists when $\chi > \pi_1$ and action 1 is optimal under the prior.[23] The following proposition generalizes this result.

Proposition 9 *Let $a = \sum_{\theta=0}^{\theta^*} p_\theta$ be the prior probability that the innovation is inefficient. Suppose there exists a set of efficient states $\Theta^* \subseteq \{\theta^*+1, \dots, n\}$ with $\sum_{\theta\in\Theta^*} p_\theta = b$ such that (i) $\chi > \pi_{\max\Theta^*}$ and (ii) action 1 is optimal when $\theta = 0$ with probability $a/(a+b)$ and $\theta = \min\Theta^*$ with probability $b/(a+b)$:*

$$\frac{a}{a+b}\pi_0 + \frac{b}{a+b}\pi_{\min\Theta^*} - c > \chi. \tag{15}$$

Then there does not exist an asymptotically efficient path.

For example, suppose that (i) there is only one inefficient state, state 0, and (ii) there is an efficient state $\hat{\theta}$ with $\chi > \pi_{\hat{\theta}}$ such that action 1 is optimal conditional on the event

[23] Conversely, Theorem 3 can be extended to show that an asymptotically efficient path does exist when $\chi < \pi_1$.

$\theta \in \left\{0, \hat{\theta}\right\}$. Then Proposition 9 implies that an asymptotically efficient path does not exist.

It seems difficult to substantially weaken the sufficient condition for the non-existence of an asymptotically efficient path in Proposition 9. However, note that the above example of an asymptotically efficient path with $n = 2$ has the curious feature that at some points in time the success rate is non-monotone in the state. Equilibrium paths with this feature are arguably implausible, because the inferences players must draw based on their samples are quite complex: for instance, in the example there are points in time where players take the risky action if they observe a small or large number of successes and take the safe action if they observe an intermediate number of successes.[24] I will thus say that an equilibrium path is *simple* if at each point in time observing success is always unambiguously good news or bad news: for each t, $x_\theta(t)(\pi_\theta - \chi)$ is either increasing or decreasing in θ.[25] Similarly, an asymptotically efficient path is simple if the corresponding equilibrium path for each K is simple. The following result gives a sufficient condition for the non-existence of a simple asymptotically efficient path exists.

Proposition 10 *If $\chi > \pi_\theta$ for some efficient state θ, then there does not exist a simple asymptotically efficient path.*

Note that, with only two states, every equilibrium path is simple. Hence, when there are only two states, Proposition 10 says precisely that there is no asymptotically efficient path in the cost-saving innovation case. Intuitively, with two states, observing success cannot switch from being good news to being bad news without passing through a point where it is completely uninformative. However, with more than two states, success can potentially switch from being good news to being bad news by passing through a region where it is "mixed news," in that it shifts the likelihood ratio of pairs of states in a non-monotone manner. As the requirement of simplicity rules out this mixed news case, results from the two-state case generalize under this restriction.

[24] In particular, this is how players behave at points in time where the success rate in state 0 has already crossed the succes rate in state 1, but has not yet crossed the success rate in state 2.

[25] Note that the success rate is increasing (resp., decreasing) in θ if and only if observing more (resp., fewer) successes is better news in the monotone likelihood ratio sense.

6.2 More General Information Structures

6.2.1 Additional Signals of the State

Theorems 1, 2, and 3 are robust to letting players observe additional exogenous signals about the state, so long as these signals are not so informative that the population shares that would result from the exogenous information alone are irregular. For example, such signals could result from (unmodeled) individual information-gathering or experimentation on the part of each player prior to making her decision. The same results apply a fortiori if players observe signals of other variables that are measurable with respect to the state, such as the aggregate adoption rate or aggregate social welfare. This is of course a crucial dimension of robustness, as realistically players will always have some information about the effectiveness of the innovation, the aggregate adoption rate, and aggregate social welfare beyond that which can be directly inferred from observing others players' outcomes.

Formally, suppose that, before a player acts, she observes an additional real-valued signal $\omega \in \mathbb{R}$ that may depend on both the state θ and the player's observed sample of K outcomes, which I denote by ζ. (Thus, $\zeta \in \{0, 1\}^K$.) Assume that ω has an atomless distribution with continuous density $f_\theta (\omega|\zeta)$ and is independent across players.

The main results continue to hold so long as the signal ω is not too informative. In particular, fixing a sample ζ, assume without loss of generality that $f_1 (\omega|\zeta) / f_0 (\omega|\zeta)$ is increasing, so higher signals are better news about the state. Define a critical signal $\omega^* (\zeta)$ by

$$\frac{f_0 (\omega^* (\zeta) |\zeta)}{f_1 (\omega^* (\zeta) |\zeta)} = \frac{p}{1-p} \frac{1-p^*}{p^*},$$

or, if no such signal exists, $\omega^* (\zeta) = -\infty$. With the above ordering on ω, let $F_\theta (\omega|\zeta) = \int_{\omega' < \omega} f_\theta (\omega'|\zeta) \, d\omega'$. The required assumption on the informativeness of ω is as follows:

Assumption 1 For every sample ζ, the point $(x_0, x_1) = (1 - F_0 (\omega^* (\zeta) |\zeta) , 1 - F_1 (\omega^* (\zeta) |\zeta))$ is regular.

Note that $F_0 (\omega^* (\zeta) |\zeta) \geq F_1 (\omega^* (\zeta) |\zeta)$. Therefore, Assumption 1 is non-trivial only in the cost-saving case ($\chi > \pi_1$), in which case it is equivalent to the inequality

$$\frac{1 - F_1 (\omega^* (\zeta) | \zeta)}{1 - F_0 (\omega^* (\zeta) | \zeta)} \leq \frac{\chi - \pi_0}{\chi - \pi_1}.$$

The left-hand side of this inequality the slope of the line from the origin to the point $(1 - F_0 (\omega^* (\zeta) | \zeta), 1 - F_1 (\omega^* (\zeta) | \zeta))$ that results when players make decisions on the basis of the exogenous signals alone, while the right-hand side is the slope of the line L bounding the regular region R. When the signal is completely uninformative, $F_0 (\omega^* (\zeta) | \zeta) = F_1 (\omega^* (\zeta) | \zeta) = 0$, so this inequality holds. Assumption 1 therefore says that the signal cannot be too informative, relative to the difference between π_0 and π_1.

It is straightforward to check that the main results hold under Assumption 1. In particular, the assumption implies that, if $x_1 (\chi - \pi_1) = x_0 (\chi - \pi_0)$ (so that other players' outcomes are uninformative), then the exogenous signals push the population dynamic toward the regular point $(1 - F_0 (\omega^* (\zeta) | \zeta), 1 - F_1 (\omega^* (\zeta) | \zeta))$. The proof of Theorem 1 is thus easily adapted to show that the population dynamic can never exit the regular region R. The proofs of Theorems 2 and 3 remain valid with trivial modifications.

The results are similarly robust to introducing "noise players," who use the innovation with some exogenously given probability, not necessarily independent of the state. As long as the shares of noise players using the innovation in the two states constitute a regular point, equilibrium paths cannot exit the regular region R. For example, if the noise players do use the innovation with the same probability in both states (so the point corresponding to their behavior lies on the 45° line in (x_0, x_1)-space) then this result applies for *any* proportion of noise players in the population. Note also that players who are simply unaware of the innovation—and thus always use the status quo—are a special case of such noise players.

6.2.2 Additional Signals of Others' Actions

One may also consider what happens if, in addition to observing K other players' outcomes, each player also observes additional signals of these players' actions.[26] This extension leads to a model that combines action-based and outcome-based learning, which can be intractable

[26]Note that this is different than observing signals of the aggregate adoption rate. As the aggregate adoption rate is measurable with respect to the state, that possibility is covered by the previous subsection.

in general. However, it is possible to assess the robustness of my results to this extension in some special cases.

Suppose first that $K = 1$, and consider the extreme case where a player can perfectly observe the action of the player whose outcome she observes.[27] In this case, Theorems 1 and 2 continue to hold, despite the fact that the improvement principle applies. To see this, note that, if $x_1 (\chi - \pi_1) = x_0 (\chi - \pi_0)$ and a player observes $(x = 1, y = 1)$, then her posterior is

$$\left[1 + \frac{1-p}{p} \frac{x_0 \pi_0}{x_1 \pi_1}\right]^{-1} > p > p^*,$$

and if she observes $(x = 1, y = 0)$ then her posterior is

$$\left[1 + \frac{1-p}{p} \frac{x_0 (1 - \pi_0)}{x_1 (1 - \pi_1)}\right]^{-1} > p > p^*,$$

as $x_1 (\chi - \pi_1) = x_0 (\chi - \pi_0)$ implies that $x_1 (1 - \pi_1) \geq x_0 (1 - \pi_0)$. In addition, a player's posterior is the same whether she observes $(x = 0, y = 1)$ or $(x = 0, y = 0)$, as the probability of success under the safe action is independent of the state. There are thus two possibilities: if the posterior after $x = 0$ is below p^*, then (x_0, x_1) itself is a regular equilibrium (as each player copies the action she observes, so population shares are constant); while if this posterior is above p^* then the population dynamic at (x_0, x_1) drifts toward the regular point $(1, 1)$. The population dynamic can therefore never exit the regular region. Furthermore, the efficiency bound (12) continues to apply at every regular point.

However, if each player perfectly observes the actions of two or more players whose outcomes she also observes, the population dynamic converges to the efficient point $(x_0, x_1) = (1, 1)$ starting from any initial point, in either the outcome-improving or cost-saving case. This follows from Theorem 1 of Banerjee and Fudenberg (2004), as the observed outcomes constitute signals which (conditional on the observed actions) "convey information" in the sense of their theorem.

In addition, if K is large and players receive independent signals of the actions of K other players, then players can almost perfectly infer the adoption rate of the innovation, and hence

[27] This case was also examined by Banerjee and Fudenberg (2004).

can back out the state whenever $x_0 \neq x_1$. (If instead $x_0 = x_1$, then $x_0(\pi_0 - \chi) \neq x_1(\pi_1 - \chi)$, so the players can infer the state on the basis of outcomes.). The result that equilibria are inefficient in the cost-saving case when K is large thus relies on players' not receiving independent signals of each other's actions. For example, in the heart disease example discussed in the introduction, this result depends on whether doctors have a large amount of data on mortality rates alone rather than large amounts of data on both mortality rates and prescription rates.

A related variation to which the main results are *not* robust is that where observed outcomes are "time-stamped," so that a player observes not only the outcomes of K players currently in the population but also the times at which they acted. In this variation, even if the current adoption rates satisfy $x_1(\pi_1 - \chi) = x_0(\pi_0 - \chi)$—so that observing a random outcome is uninformative about the state—observing time-stamped outcomes essentially lets a player observe outcomes generated by earlier adoption rates, which remain informative.

6.2.3 Unknown Calendar Time

The analysis so far assumes that, when a player enters the game, she knows the adoption rate of the innovation conditional on the state. For the definition of a steady state, this assumption is innocuous. However, as adoption rates change over time along a dynamic equilibrium path, the definition of an equilibrium path implicitly assumes that a player knows the time at which she enters the game. There are several ways of relaxing this assumption, but a particularly simple and natural one is to assume that players always draw inferences as if their samples were drawn from a fixed steady state distribution. An interpretation is that players have an improper uniform prior over the time at which they enter the game and expect a steady state to eventually be reached.

Definition 3 *Fix a steady state* (x_0^*, x_1^*, k^*, s^*). *An* unknown calendar time path relative *to* (x_0^*, x_1^*, k^*, s^*) *is a pair of differentiable functions* $(x_0 : \mathbb{R}_+ \to [0,1], x_1 : \mathbb{R}_+ \to [0,1])$ *such* *that, for* $\theta = 0, 1$*,*

$$\dot{x}_\theta(t) = \left\{ \begin{array}{c} \phi_\theta(k^*; x_\theta(t))\, s^* + \sum_{k=k^*+1}^{K} \phi_\theta(k; x_\theta(t)) - x_\theta(t) \\ \text{if (4) holds at } (x_0^*, x_1^*); \\ \phi_\theta(k^*; x_\theta(t))\, s^* + \sum_{k=0}^{k^*-1} \phi_\theta(k; x_\theta(t)) - x_\theta(t) \\ \text{if (4) fails at } (x_0^*, x_1^*) \end{array} \right\}.$$

My main results are robust to considering this alternative equilibrium concept. In particular, an unknown calendar time path relative to an irregular equilibrium can never converge to that equilibrium from a regular initial point. The intuition is that the success rate at the initial point is higher in the good state, and if players use the steady-state inference rule to mistakenly conclude that success is bad news about the state, then they switch to the safe action at a higher rate in the good state, which further improves the success rate in the good state.

Proposition 11 *If* (x_0^*, x_1^*, k^*, s^*) *is an irregular equilibrium,* (x_0, x_1) *is an unknown calendar time path relative to* (x_0^*, x_1^*, k^*, s^*)*, and* $(x_0(0), x_1(0))$ *is regular, then* $(x_0(t), x_1(t))$ *is regular for all* $t \in \mathbb{R}_+$*. In particular,* $(x_0(t), x_1(t))$ *is bounded away from* (x_0^*, x_1^*) *for all* $t \in \mathbb{R}_+$*.*

Conversely, a large class of regular steady states can be reached by an unknown calendar time path starting from a range of regular initial points, including the point $(x_0, x_1) = (1, 1)$ that results from optimal choice under the prior. Say that a regular steady state (x_0^*, x_1^*, k^*, s^*) is *stable from above* if, for $\theta = 0, 1$, $\phi_\theta(k^*; x_\theta)\, s^* + \sum_{k=k^*+1}^{K} \phi_\theta(k; x_\theta) < x_\theta$ for all $x_\theta > x_\theta^*$. When $K = 1$, the unique equilibrium is stable from above; in Example 2 in Appendix A, the larger stable regular equilibrium is stable from above.

Proposition 12 *If* (x_0^*, x_1^*, k^*, s^*) *is a regular equilibrium that is stable from above,* (x_0, x_1) *is an unknown calendar time path relative to* (x_0^*, x_1^*, k^*, s^*)*, and* $(x_0(0), x_1(0))$ *is regular and satisfies* $x_\theta(0) \geq x_\theta^*$ *for* $\theta = 0, 1$*, then* $\lim_{t\to\infty}(x_0(t), x_1(t)) = (x_0^*, x_1^*)$*.*

7 Conclusion

This paper has developed a simple model of social learning where learning is *outcome-based*: players observe each other's outcomes, but not their actions. Since outcomes are noisy, the resulting picture of social learning features both inefficiency and diversity of actions. I have focused on the question of how the nature and extent of inefficiency depend on features of the innovation about which the group must learn. The most striking finding is that, while outcome-improving innovations are adopted efficiently when either outcomes are not very noisy or players' samples are large, the adoption of cost-saving innovations entails substantial inefficiency in any equilibrium satisfying a natural regularity condition. While I have not modeled the production of innovations, this difference in adoption patterns can be expected to bias the innovation process against cost-saving innovations, which is consistent with the observed lack of cost-saving innovations in many fields.

Let me conclude by pointing out two implications of the model, one positive and one normative.

First, while the model is admittedly simple and stylized, it does make some clear empirical predictions that would be interesting to test. The key prediction is that the correlation between an environment's underlying suitability to an innovation and the realized adoption rate will be greater for outcome-improving innovations than for cost-saving innovations. For example, in the agricultural context, the prediction is that the correlation between soil suitability and adoption will be greater for outcome-improving innovations (like high-yield crop varieties) than for cost-saving innovations (like labor-saving varieties, such as the genetically engineered soybeans studied by Bustos, Caprettini, and Ponticelli (2016)). While I am not aware of existing data that could directly be brought to bear here, this seems like an interesting direction for empirical research.

Second, consider the situation of a benevolent outsider, like an NGO or professional organization, that wants to design an intervention to improve the efficiency of technology adoption. In the problematic cost-saving case, publicly releasing information (such as test results) about the technology's effectiveness may not help: the public release of information at the beginning of the game serves only to change the prior p, and the result that the equilibrium adoption of cost-saving innovations is inefficient is largely independent of the prior. But other interventions could be more effective. For instance, gathering and releasing data on the adoption rate of the innovation is a promising way to restore efficiency: even when equilibrium adoption is inefficient, the adoption rate will still differ depending on the innovation's effectiveness, so revealing the adoption rate will reveal the effectiveness of the innovation. Another promising approach is providing an information technology that lets each individual conduct her own test of the innovation before deciding whether to adopt it. If these tests are sufficiently informative, the initial adoption rates will be irregular (i.e., Assumption 1 will be violated), and social learning can then lead the adoption rates to converge to efficiency.

8 Appendix A: Supporting Results and Examples

This appendix contains results supporting the analysis of Section 3 and 4, as well as examples of irregular equilibria.

A preliminary result records two simple but useful properties of equilibria: in any equilibrium (regular or irregular), the share of the population experimenting is higher in the good state, and the share experimenting in the good state is bounded away from 0.

Proposition 13 *In any equilibrium,* $x_1 \geq x_0$ *and* $x_1 \geq \frac{p-p^*}{1-p^*}$.

It immediately follows that an irregular equilibrium can exist only in the case of a cost-saving innovation.

Proposition 14 *If* $\chi \leq \pi_1$*, then every equilibrium is regular.*

Proof. If $\pi_0 \leq \chi \leq \pi_1$, this follows because (4) holds for all $x_0, x_1 \geq 0$.

If $\chi < \pi_0$, then (4) holds whenever $x_1 \geq x_0$, so the result follows from Proposition 13. ■

The next result is that there is a unique equilibrium when $K = 1$. By Theorem 1, it must be regular. The proof fully characterizes the equilibrium.

Proposition 15 *When* $K = 1$*, there is a unique equilibrium. In the equilibrium, players experiment with probability 1 after observing a success and experiment with probability less than 1 after observing a failure.*

Irregular equilibria can however exist in the cost-saving innovation case when $K > 1$, as the following examples show. The second example also shows that a *stable* irregular equilibrium can coexist with a stable regular equilibrium (and indeed with more than one of them).

Example 1: An Irregular Equilibrium

Let $K = 2$, $\chi = 1$, $\pi_0 = 0$, $\pi_1 = \frac{1}{3}$, $p = \frac{1}{2}$, and $c = -\frac{8}{9}$. I claim that the irregular pair $\left(x_0 = 0, x_1 = \frac{3}{4}\right)$, together with the strategy of experimenting if and only if at least one observation is a failure, is an equilibrium. (The pair is irregular because the success rate is 1 in state 0 and $1 - x_1 (1 - \pi_1) = \frac{1}{2}$ in state 1.)

This follows because $p^* = \frac{\chi + c - \pi_0}{\pi_1 - \pi_0} = \frac{1}{3}$, while the posterior probability that $\theta = 1$ after observing at least one failure is $1 > p^*$, and the posterior probability that $\theta = 1$ after observing zero failures is

$$\left[1 + \frac{1-p}{p}\frac{1}{\left(\frac{1}{2}\right)^2}\right]^{-1} = \frac{1}{5} < p^*.$$

The stated strategy is therefore optimal. In addition, (x_0, x_1) is an stationary point because the probability of observing at least one failure is 0 in state 0 and $1 - \left(\frac{1}{2}\right)^2 = \frac{3}{4}$ in state 1.

The equilibrium is however unstable, as the probability of observing at least one failure in state 0 when fraction x_0 experiments equals $1 - (1 - x_0)^2$, which is greater than x_0 for all $x_0 \in (0, 1)$.

Example 2: A Stable Irregular Equilibrium (and Two Stable Regular Equilibria)

Let $K = 3$, $\chi = \frac{9}{10}$, $\pi_0 = 0$, $\pi_1 = \frac{1}{10}$, $p = \frac{1}{2}$, and $c = -\frac{1701}{2000}$.[28] Under the strategy of experimenting if and only if at least two observations are failures, the equation for x_θ to be a stationary point is

$$x_\theta = (1 - \chi + x_\theta (\chi - \pi_\theta))^3 + 3(1 - \chi + x_\theta (\chi - \pi_\theta))^2 (\chi - x_\theta (\chi - \pi_\theta)).$$

Consider the pair (x_0, x_1) given by taking the smallest solution to this cubic equation for $\theta = 0$ and the largest solution for $\theta = 1$: $(x_0, x_1) \approx (.07407, .9419)$. This pair is irregular because the success rate is $\chi - x_1 (\chi - \pi_1) \approx 0.1465$ in state 1 and $\chi - x_0 (\chi - \pi_0) \approx 0.8333$ in state 0. It is straightforward to check that this pair is stable: for $\theta = 0, 1$, the above cubic equation has three roots, of which the middle one is unstable. Finally, to see that the proposed strategy is optimal, note that $p^* = \frac{\chi + c - \pi_0}{\pi_1 - \pi_0} = 0.495$, while the posterior probability that $\theta = 1$ after observing two failures is

$$\left[1 + \frac{1-p}{p}\frac{(1 - \chi + x_0 (\chi - \pi_0))^2 (\chi - x_0 (\chi - \pi_0))}{(1 - \chi + x_1 (\chi - \pi_1))^2 (\chi - x_1 (\chi - \pi_1))}\right]^{-1} \approx 0.8217 > p^*,$$

[28] The explanation for this oddly precise choice of c is that, if $c = -\frac{17}{20}$, the analysis of the example would be exactly the same except that one would have $p^* = p$, which violates (1). As the only role of c in the model is to determine p^*, it suffices to let $c = -\frac{17}{20} - \varepsilon$ for any sufficiently small $\varepsilon > 0$.

while the posterior probability that $\theta = 1$ after observing one failure is

$$\left[1+\frac{1-p}{p}\frac{\left(1-\chi+x_0\left(\chi-\pi_0\right)\right)\left(\chi-x_0\left(\chi-\pi_0\right)\right)^2}{\left(1-\chi+x_1\left(\chi-\pi_1\right)\right)\left(\chi-x_1\left(\chi-\pi_1\right)\right)^2}\right]^{-1}\approx 0.1366 < p^*.$$

In fact, it is not hard to see that a stable irregular equilibrium cannot exist when $K = 2$, so $K = 3$ is the minimum sample size for which a stable irregular equilibrium can exist. The reason is that, when $K = 2$, the fraction of players observing at least k failures in state θ is at most quadratic in x_θ, so there is a unique stable irregular stationary point (x_0, x_1). But, since failure is more likely for a given fraction of experimenters in state 0, this unique stable pair always has $x_0 > x_1$, so by Proposition 13 it cannot be an equilibrium.

This same example also admits two stable regular equilibria. Thus, there can be multiple stable regular equilibria, and they can coexist with a stable irregular equilibrium.

Specifically, I claim that a pair $(x_0', x_1') \approx (0.4681, 0.5)$, together with the strategy of experimenting if and only if at least two successes are observed, is an equilibrium; and that so is a pair $(x_0'', x_1'') \approx (0.6625, 0.7061)$, together with the strategy of experimenting if and only if at least one success is observed. The intuition for this multiplicity is that, when the "bar" for experimenting is raised from one observed success to two, this reduces the steady-state experimentation rate, which makes failure less likely in both states (as $\chi > \pi_0, \pi_1$), and thus makes failure more informative. This in turn justifies the greater number of observed successes required for experimentation. The formal construction of the two regular equilibria is deferred to Appendix B.

Finally if one considers this example with $K = 2$ rather than $K = 3$, one finds that there is a unique stable regular equilibrium $(x_0, x_1) \approx (0.5955, 0.6327)$ (corresponding to the strategy of experimenting if and only if at least one success if observed), and welfare in this steady state lies in between that in the two stable regular steady states that arise when $K = 3$. This shows that welfare does not always unambiguously increase when players observe larger samples, even within the class of stable regular equilibria.

References

[1] Acemoglu, Daron, and Alexander Wolitzky (2014), "Cycles of Conflict: An Economic Model," *American Economic Review*, 104: 1350-1367.

[2] Aghion, Philippe, Patrick Bolton, Christopher Harris, and Bruno Jullien (1991), "Optimal Learning by Experimentation," *Review of Economic Studies*, 58: 621-654.

[3] Aubin, J.-P., and A. Cellina (1984), *Differential Inclusions: Set-Valued Maps and Viability Theory*, Springer: Berlin.

[4] Bala, Venkatesh, and Sanjeev Goyal (1995), "A Theory of Learning with Heterogeneous Agents," *International Economic Review*, 36: 303-323.

[5] Banerjee, Abhijit V. (1992), "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107: 797-817.

[6] Banerjee, Abhijit V. (1993), "The Economics of Rumours," *Review of Economic Studies*, 60: 309-327.

[7] Banerjee, Abhijit, and Drew Fudenberg (2004), "Word-of-Mouth Learning," *Games and Economic Behavior*, 46: 1-22.

[8] Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992), "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100: 992-1026.

[9] Bloom, Nicholas, Ben Eifert, Aprajit Mahajan, David McKenzie, and John Roberts (2013), "Does Management Matter? Evidence from India," *Quarterly Journal of Economics*, 128: 1-51.

[10] Bloom, Nicholas, and John Van Reenen (2010), "Why Do Management Practices Differ across Firms and Countries?" *Journal of Economic Perspectives*, 24: 203-224.

[11] Bolton, Patrick, and Christopher Harris (1999), "Strategic Experimentation," *Econometrica*, 67: 349-374.

[12] Bustos, Paula, Bruno Caprettini, and Jacopo Ponticelli (2016), "Agricultural Productivity and Structural Transformation: Evidence from Brazil," *American Economic Review*, 106: 1320-1365.

[13] Chakma, Justin, Gordon H. Sun, Jeffrey D. Steinberg, Stephen M. Sammut, and Reshma Jagsi (2014), "Asia's Ascent—Global Trends in Biomedical R&D Expenditures," *New England Journal of Medicine*, 370: 3-6.

[14] Chamley, Christophe (2004), *Rational Herds: Economic Models of Social Learning*, Cambridge University Press: Cambridge.

[15] Chandra, Amitabh, and Jonathan Skinner (2012), "Technology Growth and Expenditure Growth in Health Care," *Journal of Economic Literature*, 50: 645-680.

[16] Che, Yeon-Koo, and Johannes Hörner (2016), "Optimal Design for Social Learning," *working paper*.

[17] Conley, Timothy G., and Christopher R. Udry (2010), "Learning About a New Technology: Pineapple in Ghana," *American Economic Review*, 100: 35-69.

[18] Cutler, David M. (2005), *Your Money or Your Life: Strong Medicine for America's Health Care System*, Oxford University Press: Oxford.

[19] Das, Jishnu, Alaka Holla, Aakash Mohpol, and Karthik Muralidharan, (2016) "Quality and Accountability in Healthcare Delivery: Audit-Study Evidence from Primary Care in India," *American Economic Review*, Forthcoming.

[20] Duflo, Esther, Michael Kremer, and Jonathan Robinson (2008), "How High are Rates of Return to Fertilizer? Evidence from Field Experiments in Kenya," *American Economic Review: Papers & Proceedings*, 98: 482-488.

[21] Easley, David, and Nicholas M. Kiefer (1988), "Controlling a Stochastic Process with Unknown Parameters," *Econometrica*, 56: 1045-1064.

[22] Foster, Andrew D., and Mark R. Rosenzweig (2010), "Microeconomics of Technology Adoption," *Annual Review of Economics*, 2: 395-424.

[23] Frick, Mira and Yuhta Ishii (2016), "Innovation Adoption by Forward-Looking Social Learners," *working paper*.

[24] Hörner, Johannes and Andrzej Skrzypacz (2016), "Learning, Experimentation and Information Design," *working paper*.

[25] Keller, Godfrey, Sven Rady, and Martin Cripps (2005), "Strategic Experimentation with Exponential Bandits," *Econometrica*, 73: 39-68.

[26] Kiefer, Nicholas M. (1989), "A Value Function Arising in the Economics of Information," *Journal of Economic Dynamics and Control*, 13: 201-223.

[27] Kremer, Ilan, Yishay Mansour, and Motty Perry (2014), "Implementing the "Wisdom of the Crowd"," *Journal of Political Economy*, 122: 988-1012.

[28] Lauermann, Stephan, and Asher Wolinsky (2016), "Search with Adverse Selection," *Econometrica*, 84: 243-315.

[29] McLennan, Andrew (1984), "Price Dispersion and Incomplete Learning in the Long Run," *Journal of Economic Dynamics and Control*, 7: 331-347.

[30] Milgrom, Paul R. (1981), "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 12: 380-391.

[31] Murto, Pauli, and Juuso Välimäki (2011), "Learning and Information Aggregation in an Exit Game," *Review of Economic Studies*, 78: 1426-1461.

[32] Nelson, Aaron L., Joshua T. Choen, Dan Greenberg, and David M. Kent, (2009), "Much Cheaper, Almost as Good: Decrementally Cost-Effective Medical Innovation," *Annals of Internal Medicine*, 151: 662-667.

[33] Newhouse, Joseph P. (1992), "Medical Care Costs: How Much Welfare Loss?" *Journal of Economic Perspectives*, 6: 3-21.

[34] Piketty, Thomas (1995), "Social Mobility and Redistributive Politics," *Quarterly Journal of Economics*, 110: 551-584.

[35] Pray, Carl E., and Keith O. Fuglie, (2015), "Agricultural Research by the Private Sector," *Annual Review of Resource Economics*, 7: 399-424.

[36] Schlibach, Frank (2015), *Essays in Development and Behavioral Economics*, PhD Thesis, Harvard University.

[37] Smith, Lones, and Peter Sørensen (2000), "Pathological Outcomes of Observational Learning," *Econometrica*, 68: 371-398.

[38] Smith, Lones, and Peter Sørensen (2014), "Rational Social Learning with Random Sampling," *working paper.*

[39] Wagner, Peter (2017), "Who Goes First? Strategic Delay under Information Asymmetry," *Theoretical Economics*, forthcoming.

9 Appendix B: Omitted Proofs (For Online Publication)

In what follows, the results of Appendix A are proved before those of Sections 3, 4, and 5.

9.1 Properties of F

Non-empty: If (x_0, x_1) satisfies (4) (resp., the opposite of (4)) then $p(k; x_0, x_1)$ is increasing (resp., decreasing) in k. In either case, $p(k; x_0, x_1)$ is monotone in k, so there exists $k^* \in \{0, \dots, K\}$ such that either $p(k^* - 1; x_0, x_1) \leq p^* \leq p(k^* + 1; x_0, x_1)$ or $p(k^* - 1; x_0, x_1) \geq p^* \geq p(k^* + 1; x_0, x_1)$, where in both cases the first (last) inequality is vacuous if $k^* = 0$ ($k^* = K$). With this value of k^*, let $s = 1$ if $p(k^*; x_0, x_1) \geq p^*$ and let $s = 0$ if $p(k^*; x_0, x_1) < p^*$. Next, with these values of k^* and s, let (x_0', x_1') be computed as in Definition 1.

Convex-valued: Recall that there is at most one value of $k^* \in \{0, \dots, K\}$ such that $p(k^*; x_0, x_1) = p^*$. So, if there are distinct elements of $F(x_0, x_1)$, (x_0', x_1') and (x_0'', x_1''), it must be that (x_0', x_1') and (x_0'', x_1'') are computed as in (8) or (9) with distinct values $s', s'' \in [0, 1]$. But then, for all $\alpha \in [0, 1]$, letting $s = \alpha s' + (1 - \alpha) s''$, it follows that $(\alpha x_0' + (1 - \alpha) x_0'', \alpha x_1' + (1 - \alpha) x_1'') \in F(x_0, x_1)$.

Compact-valued: Fix (x_0, x_1) and a sequence $(x_0'^n, x_1'^n) \to (x_0', x_1')$, with $(x_0'^n, x_1'^n) \in F(x_0, x_1)$, and let $(k^{*,n}, s^n)$ be arbitrarily chosen corresponding values of k^* and s. Taking any convergent subsequence of $(k^{*,n}, s^n) \to (k^{*,\infty}, s^\infty)$, it follows from continuity of ϕ_θ that, with $k^* = k^{*,\infty}$ and $s = s^\infty$, (x_0', x_1') satisfy the conditions for inclusion in $F(x_0, x_1)$.

Upper hemi-continuous: Fix sequences $(x_0^n, x_1^n) \to (x_0, x_1)$ and $(x_0'^n, x_1'^n) \to (x_0', x_1')$, with $(x_0'^n, x_1'^n) \in F(x_0^n, x_1^n)$, and let $(k^{*,n}, s^n)$ be arbitrarily chosen corresponding values of k^* and s. Taking any convergent subsequence of $(k^{*,n}, s^n) \to (k^{*,\infty}, s^\infty)$, it follows from continuity of ϕ_θ that, with $k^* = k^{*,\infty}$ and $s = s^\infty$, (x_0', x_1') satisfy the conditions for inclusion in $F(x_0, x_1)$. ■

9.2 Omitted Step in the Proof of Theorem 1

To show that an equilibrium path with a regular initial point visits only regular points, note that, by assumption, $x_0(0)(\pi_0 - \chi) \leq x_1(0)(\pi_1 - \chi)$. As x_0 and x_1 are continuous, if there exists a time t' with $x_0(t')(\pi_0 - \chi) > x_1(t')(\pi_1 - \chi)$, then by the intermediate value theorem and the definition of the derivative there must exist another time t where $x_0(t)(\pi_0 - \chi) = x_1(t)(\pi_1 - \chi)$ but it is not the case that x_0 and x_1 are differentiable at t with $\dot{x}_0(t)(\pi_0 - \chi) < \dot{x}_1(t)(\pi_1 - \chi)$ (in particular, the time $\sup\{t < t' : x_0(t)(\pi_0 - \chi) = x_1(t)(\pi_1 - \chi)\}$ must have this property). But, if $x_0(t)(\pi_0 - \chi) = x_1(t)(\pi_1 - \chi)$ then $p(k; x_0(\tau), x_1(\tau)) = p > p^*$

for all $k \in \{0, \ldots, K\}$ and all τ in a neighborhood of t, and hence x_θ is differentiable at t with $\dot{x}_\theta(t) = 1 - x_\theta(t)$ for $\theta = 0, 1$.[29] Therefore,

$$\begin{aligned}
\dot{x}_0(t)(\pi_0 - \chi) &= \pi_0 - \chi - x_0(t)(\pi_0 - \chi) \\
&= \pi_0 - \chi - x_1(t)(\pi_1 - \chi) \\
&< \pi_1 - \chi - x_1(t)(\pi_1 - \chi) = \dot{x}_1(t)(\pi_1 - \chi).
\end{aligned}$$

It follows that there can be no such time t'. ■

9.3 Proof of Proposition 13

In a regular equilibrium (x_0, x_1, k^*, s), $\phi_0(k; x_0)/\phi_1(k; x_1)$ is decreasing in k, and therefore (as likelihood ratio dominance implies first-order stochastic dominance)

$$\sum_{k=k_0}^{K} \phi_0(k; x_0) \leq \sum_{k=k_0}^{K} \phi_1(k; x_0) \text{ for all } k_0 \in \{0, \ldots, K\}.$$

In particular, this inequality holds for $k_0 = k^*$ and $k_0 = k^* + 1$, so

$$x_0 = \phi_0(k^*; x_0)s + \sum_{k=k^*+1}^{K} \phi_0(k; x_0) \leq \phi_1(k^*; x_1)s + \sum_{k=k^*+1}^{K} \phi_1(k; x_1) = x_1.$$

Symmetrically, in an irregular equilibrium, $\phi_0(k; x_0)/\phi_1(k; x_1)$ is increasing in k and players experiment if $k < k^*$, so again $x_0 \leq x_1$.

The second part of the proposition follows because the Bayesian information structure that minimizes the probability that a player's posterior is strictly above p^* sends her posterior to p^* and 1 with probabilities $\frac{1-p}{1-p^*}$ and $\frac{p-p^*}{1-p^*}$, respectively. ■

[29] Proof: If $p(k; x_0(\tau), x_1(\tau)) > p^*$ for all $k \in \{0, \ldots, K\}$ and all τ in a neighborhood of t, then $\dot{x}_\theta(\tau) = 1 - x_\theta(\tau)$ for all τ at which x_θ is differentiable in a neighborhood of t. But then

$$\lim_{\varepsilon \to 0} \frac{1}{\varepsilon}(x_\theta(t+\varepsilon) - x_\theta(t)) = \lim_{\varepsilon \to 0} \frac{1}{\varepsilon} \int_t^{t+\varepsilon} \dot{x}_\theta(\tau)\, d\tau = \lim_{\varepsilon \to 0} \frac{1}{\varepsilon} \int_t^{t+\varepsilon} (1 - x_\theta(\tau))\, d\tau = 1 - x_\theta(t),$$

where the first equality uses absolute continuity of x and the last uses continuity.

9.4 Proof of Proposition 15

Fix an equilibrium, and suppose players experiment with probability s_1 after observing a success and experiment with probability s_0 after observing a failure. Then, for $\theta = 0, 1$,

$$x_\theta = [\chi + x_\theta (\pi_\theta - \chi)] s_1 + [1 - \chi - x_\theta (\pi_\theta - \chi)] s_0,$$

or

$$x_\theta = \frac{s_0 + \chi (s_1 - s_0)}{1 - (\pi_\theta - \chi)(s_1 - s_0)}. \tag{16}$$

Suppose toward a contradiction that $s_0 = 1$. As $s_0 = s_1 = 1$ would lead to $x_0 = x_1 = 1$, which is not an equilibrium by (2), this implies that $s_0 > s_1$. But $s_0 > s_1$ implies that $x_0 > x_1$, which contradicts Proposition 13. Hence, $s_0 < 1$.

Now, $s_0 < 1$ implies that $p(0; x_0, x_1) \leq p^*$. As $p > p^*$ and p is a convex combination of $p(0; x_0, x_1)$ and $p(1; x_0, x_1)$ (by the law of total probability), this implies that $p(1; x_0, x_1) > p^*$. Hence, $s_1 = 1$.

Next, using (16) and $s_1 = 1$,

$$p(0; x_0, x_1) = \left[1 + \frac{1 - \pi_0}{1 - \pi_1} \frac{1 - (\pi_1 - \chi)(1 - s_0)}{1 - (\pi_0 - \chi)(1 - s_0)} \frac{1 - p}{p}\right]^{-1}.$$

Therefore, $(s_0 = 0, s_1 = 1)$ corresponds to an equilibrium if and only if

$$\frac{1 - \pi_0}{1 - \pi_1} \frac{1 - \pi_1 + \chi}{1 - \pi_0 + \chi} \geq \frac{p}{1 - p} \frac{1 - p^*}{p^*}. \tag{17}$$

On the other hand, $(s_0 = s, s_1 = 1)$ with $s > 0$ corresponds to an equilibrium if and only if

$$\frac{1 - \pi_0}{1 - \pi_1} \frac{1 - (\pi_1 - \chi)(1 - s)}{1 - (\pi_0 - \chi)(1 - s)} = \frac{p}{1 - p} \frac{1 - p^*}{p^*}. \tag{18}$$

The left-hand side of (18) is increasing in s, and by (2) it exceeds the right-hand side when $s = 1$. Hence, by the intermediate value theorem, either there is a unique equilibrium given by $(s_0 = 0, s_1 = 1)$ and (16), or there exists a unique value $s > 0$ such that the unique equilibrium is given by $(s_0 = s, s_1 = 1)$, (18), and (16). ■

9.5 Regular Equilibria in Example 2

In what follows, all parameters are as in Example 2.

Under the strategy of experimenting if and only if at least two successes are observed, the equation for x_θ to be a stationary point is

$$x_\theta = \left(\chi - x_\theta\left(\chi - \pi_\theta\right)\right)^3 + 3\left(\chi - x_\theta\left(\chi - \pi_\theta\right)\right)^2\left(1 - \chi + x_\theta\left(\chi - \pi_\theta\right)\right).$$

Let (x_0', x_1') be the unique solutions to this equation for $\theta = 0, 1$, given by $(x_0', x_1') \approx (0.4681, 0.5)$. Then the posterior probability that $\theta = 1$ after observing two successes is

$$\left[1 + \frac{1-p}{p}\frac{\left(1 - \chi + x_0'\left(\chi - \pi_0\right)\right)\left(\chi - x_0'\left(\chi - \pi_0\right)\right)^2}{\left(1 - \chi + x_1'\left(\chi - \pi_1\right)\right)\left(\chi - x_1'\left(\chi - \pi_1\right)\right)^2}\right]^{-1} \approx 0.5113 > p^*,$$

while the posterior probability that $\theta = 1$ after observing one success is

$$\left[1 + \frac{1-p}{p}\frac{\left(1 - \chi + x_0'\left(\chi - \pi_0\right)\right)^2\left(\chi - x_0'\left(\chi - \pi_0\right)\right)}{\left(1 - \chi + x_1'\left(\chi - \pi_1\right)\right)^2\left(\chi - x_1'\left(\chi - \pi_1\right)\right)}\right]^{-1} \approx 0.4900 < p^*.$$

So this is an equilibrium. It is also easily seen to be stable, as the curve $\left(\chi - x_\theta\left(\chi - \pi_\theta\right)\right)^3 + 3\left(\chi - x_\theta\left(\chi - \pi_\theta\right)\right)^2\left(1 - \chi + x_\theta\left(\chi - \pi_\theta\right)\right)$ crosses x_θ from above, for $\theta = 0, 1$.

Similarly, under the strategy of experimenting if and only if at least one successes is observed, the equation for x_θ to be a stationary point is given by

$$x_\theta = 1 - \left(1 - \chi + x_\theta\left(\chi - \pi_\theta\right)\right)^3.$$

Let (x_0'', x_1'') be the unique solutions, given by $(x_0'', x_1'') \approx (0.6625, 0.7061)$. Then the posterior probability that $\theta = 1$ after observing one successes is

$$\left[1 + \frac{1-p}{p}\frac{\left(1 - \chi + x_0''\left(\chi - \pi_0\right)\right)^2\left(\chi - x_0''\left(\chi - \pi_0\right)\right)}{\left(1 - \chi + x_1''\left(\chi - \pi_1\right)\right)^2\left(\chi - x_1''\left(\chi - \pi_1\right)\right)}\right]^{-1} \approx 0.5015 > p^*,$$

while the posterior probability that $\theta = 1$ after observing zero successes is

$$\left[1 + \frac{1-p}{p}\frac{\left(1 - \chi + x_0''\left(\chi - \pi_0\right)\right)^3}{\left(1 - \chi + x_1''\left(\chi - \pi_1\right)\right)^3}\right]^{-1} \approx 0.4655 < p^*.$$

So this is also an equilibrium, and it is also easily seen to be stable.

9.6 Proof of Theorem 3

It suffices to prove the theorem for sequences $(x_{0,K}, x_{1,K}, k_K^*, s_K)$ where $(x_{0,K}, x_{1,K})$ converges. Fix such a sequence, and let $(x_0, x_1) = \lim_{K\to\infty}(x_{0,K}, x_{1,K})$. As $p\left(k_K^* - 1; x_0^K, x_{1,K}\right) \leq p^* \leq$

$p\left(k_K^*+1; x_{0,K}, x_{1,K}\right)$, it follows that

$$\begin{aligned}
&\frac{\left(\chi+x_{0,K}\left(\pi_0-\chi\right)\right)^{k-1}\left(1-\chi-x_{0,K}\left(\pi_0-\chi\right)\right)^{K-k+1}}{\left(\chi+x_{1,K}\left(\pi_1-\chi\right)\right)^{k-1}\left(1-\chi-x_{1,K}\left(\pi_1-\chi\right)\right)^{K-k+1}} \\
\geq\ & \frac{p}{1-p}\frac{1-p^*}{p^*} \\
\geq\ & \frac{\left(\chi+x_{0,K}\left(\pi_0-\chi\right)\right)^{k+1}\left(1-\chi-x_{0,K}\left(\pi_0-\chi\right)\right)^{K-k-1}}{\left(\chi+x_{1,K}\left(\pi_1-\chi\right)\right)^{k+1}\left(1-\chi-x_{1,K}\left(\pi_1-\chi\right)\right)^{K-k-1}}.
\end{aligned}$$

Letting $\kappa_K=\frac{k_K^*}{K}$, rewrite this as

$$\begin{aligned}
&\left[\frac{\left(\chi+x_{0,K}\left(\pi_0-\chi\right)\right)^{\kappa_K-\frac{1}{K}}\left(1-\chi-x_{0,K}\left(\pi_0-\chi\right)\right)^{1-\kappa_K+\frac{1}{K}}}{\left(\chi+x_{1,K}\left(\pi_1-\chi\right)\right)^{\kappa_K-\frac{1}{K}}\left(1-\chi-x_{1,K}\left(\pi_1-\chi\right)\right)^{1-\kappa_K+\frac{1}{K}}}\right]^K \\
\geq\ & \frac{p}{1-p}\frac{1-p^*}{p^*} \\
\geq\ & \left[\frac{\left(\chi+x_{0,K}\left(\pi_0-\chi\right)\right)^{\kappa_K+\frac{1}{K}}\left(1-\chi-x_{0,K}\left(\pi_0-\chi\right)\right)^{1-\kappa_K-\frac{1}{K}}}{\left(\chi+x_{1,K}\left(\pi_1-\chi\right)\right)^{\kappa_K+\frac{1}{K}}\left(1-\chi-x_{1,K}\left(\pi_1-\chi\right)\right)^{1-\kappa_K-\frac{1}{K}}}\right]^K.
\end{aligned}$$

As this holds for all K, it follows that

$$\lim_{K\to\infty}\frac{\left(\chi+x_{0,K}\left(\pi_0-\chi\right)\right)^{\kappa_K}\left(1-\chi-x_{0,K}\left(\pi_0-\chi\right)\right)^{1-\kappa_K}}{\left(\chi+x_{1,K}\left(\pi_1-\chi\right)\right)^{\kappa_K}\left(1-\chi-x_{1,K}\left(\pi_1-\chi\right)\right)^{1-\kappa_K}}=1.$$

Now, $x_0\left(\pi_0-\chi\right)<x_1\left(\pi_1-\chi\right)$ by (13), so

$$\lim_{K\to\infty}\kappa_K=\left[1+\frac{\log\frac{\chi+x_0(\pi_0-\chi)}{\chi+x_1(\pi_1-\chi)}}{\log\frac{1-\chi-x_1(\pi_1-\chi)}{1-\chi-x_0(\pi_0-\chi)}}\right]^{-1}\in\left(\chi+x_0\left(\pi_0-\chi\right),\chi+x_1\left(\pi_1-\chi\right)\right).$$

Finally, by the weak law of large numbers for triangular arrays, as $K\to\infty$ the fraction of successes a player observes in state θ, when the probability that each observation is a success is $\chi+x_{\theta,K}\left(\pi_\theta-\chi\right)$, converges in probability to $\chi+x_\theta\left(\pi_\theta-\chi\right)$. Therefore, the probability that $k<k^*$ converges to 1 in state 0 and the probability that $k>k^*$ converges to 1 in state 1. Hence, $\lim_{K\to\infty}x_{0,K}=0$ and $\lim_{K\to\infty}x_{1,K}=1$. ■

9.7 Proof of Proposition 4

1. If $\chi>0$, then (17) holds when π_1 is close enough to 1. In this case, (16) gives $x_\theta=\frac{\chi}{1-\pi_\theta+\chi}$ for $\theta=0,1$. Hence, $x_0=\frac{\chi}{1-\pi_0+\chi}$ and $x_1\to 1$ as $\pi_1\to 1$.

2. If $\chi=0$, then (17) is violated (as $\frac{p}{1-p}\frac{1-p^*}{p^*}>1$), so the unique equilibrium is given by

(18) and (16). Solving for x_0, x_1, and s gives

$$\begin{aligned} x_0 &= \frac{(p-p^*)(1-\pi_1)}{p^*(1-p)(\pi_1-\pi_0)}, \\ x_1 &= \frac{(p-p^*)(1-\pi_0)}{p(1-p^*)(\pi_1-\pi_0)}, \text{ and} \\ s &= \frac{(p-p^*)(1-\pi_0)(1-\pi_1)}{p^*(1-p)(1-\pi_0)\pi_1 - p(1-p^*)\pi_0(1-\pi_1)}. \end{aligned}$$

Noting that $p^* \to \hat{p}$ as $\pi_1 \to 1$, it follows that $x_0 \to 0$ and $x_1 \to \frac{p-\hat{p}}{p(1-\hat{p})}$ as $\pi_1 \to 1$. ■

9.8 Proof of Proposition 5

If (17) fails, players experiment after observing success and mix after observing failure. As players must be indifferent after observing failure, welfare is the same as if players always experiment. Hence, welfare equals $p\pi_1 + (1-p)\pi_0 - c$, which is increasing in π_0, π_1, and p, decreasing in c, and constant in χ.

If (17) holds, then welfare in state θ equals

$$\chi + x_\theta(\pi_\theta - c - \chi) = \chi + \frac{\chi}{1-\pi_\theta+\chi}(\pi_\theta - c - \chi) = \frac{\chi}{1-\pi_\theta+\chi}(1-c).$$

This expression is increasing in π_θ, decreasing in c, and non-decreasing in χ, and does not depend on $\pi_{1-\theta}$. It is also greater when $\theta = 1$ than when $\theta = 0$. Hence, expected welfare is increasing in π_0, π_1 and p, decreasing in c, and non-decreasing in χ.

Finally, the mixing probability s defined by (18) is continuous in the parameters of the model and equals 0 when (17) holds with equality. Expected welfare is therefore continuous when passing from the region of parameter space where (17) holds to the region where (17) fails. The fact that expected welfare obeys the desired comparative statics within each region thus implies that the comparative statics hold globally. ■

9.9 Proof of Proposition 6

1. Fix a sequence of parameters $(\pi_0^n, \pi_1^n) \to (\pi_0, \pi_1) = (0,1)$ and fix a corresponding sequence of equilibria $(x_0^n, x_1^n, k^{*,n}, s^n) \to (x_0, x_1, k^*, s)$. Note that (x_0, x_1, k^*, s) must be an equilibrium. Suppose toward a contradiction that $x_1 < 1$. By Proposition 13, $x_1 \geq \frac{p-p^*}{1-p^*}$, so $x_1\pi_1 \in (0,1)$. On the other hand, $x_0\pi_0 = 0$. Therefore, $p(k; x_0, x_1) > p^*$ for all $k \geq 1$, and hence $k^* = 0$. The steady state equation then implies that

$$x_1 = 1 - (1-x_1)^K (1-s) \geq 1 - (1-x_1)^2 = x_1(2-x_1).$$

But this implies that $x_1 = 1$, a contradiction.

To show that $x_0 = 0$, let $\hat{s}$ be the probability with which players experiment after observing K failures in the equilibrium (x_0, x_1, k^*, s). (Thus, $\hat{s} = 0$ if $k^* > 0$, and $\hat{s} = s$ if $k^* = 0$.) Then the steady state equation implies that $x_0 = \hat{s}$. Next, note that $p(0; x_0, 1) = p(0; 1, 1) < p^*$ (by $\pi_0 = 0$ and (2)). Hence, $\hat{s} = 0$.

2. Fix a sequence of parameters $(\pi_0^n, \pi_1^n) \to (\pi_0, \pi_1) = (0, 1)$ and a corresponding sequence of regular equilibria $(x_0^n, x_1^n, k^{*,n}, s^n) \to (x_0, x_1, k^*, s)$. Note that (x_0, x_1, k^*, s) must be a regular equilibrium. I claim that $x_0 > 0$. To see this, note that, in any regular equilibrium $p(K; x_0, x_1) > p > p^*$, and therefore players experiment with probability 1 after observing K successes. Thus, if $x_0 = 0$ and $\chi = 1$, then in state 0 players would observe all successes with probability 1; and therefore x_0 would equal 1, a contradiction.

Next, as $x_0 > 0$, $\pi_0 = 0$, and $\chi + x_1(\pi_1 - \chi) = 1$, $p(k; x_0, x_1) = 0$ for all $k < K$, so players experiment with probability 0 after observing even a single failure. On the other hand, I have shown that players experiment with probability 1 after observing all successes, so

$$x_\theta = (1 - x_\theta(1 - \pi_\theta))^K \text{ for } \theta = 0, 1.$$

As $\pi_0 = 0$ and $\pi_1 = 1$, this implies that $x_0 = (1 - x_0)^K$ and $x_1 = 1$.

The last part of the proposition follows as the solution to the equation $x_0 = (1 - x_0)^K$ converges to 0 as $K \to \infty$. ■

9.10 Proof of Proposition 7

At state (x_0, x_1), the posterior belief that $\theta = 1$ after observing failure equals

$$\left[1 + \frac{1-p}{p}\frac{1 - \chi - x_0(\pi_0 - \chi)}{1 - \chi - x_1(\pi_1 - \chi)}\right]^{-1}.$$

This posterior equals p^* if and only if

$$\frac{1 - \chi - x_0(\pi_0 - \chi)}{1 - \chi - x_1(\pi_1 - \chi)} = \frac{p}{1-p}\frac{1-p^*}{p^*}.$$

This equation defines a line $\hat{L}$ in (x_0, x_1) space. Let H be half-space where the posterior exceeds p^* and let H^C be the half-space where the posterior is less than p^*; thus $\hat{L}$ marks the boundary between H and H^C. Recall from the proof of Proposition 15 that there are two possible cases: either the equilibrium is $\left(x_0 = \frac{\chi}{1-\pi_0+\chi}, x_1 = \frac{\chi}{1-\pi_1+\chi}\right)$ and this point lies in the half-space H^C, or the equilibrium lies on the line $\hat{L}$.

At any point $(x_0, x_1) \in H$, it follows that $\dot{x}_\theta = 1 - x_\theta$ for $\theta = 0, 1$, so the vector $(\dot{x}_0, \dot{x}_1)$

points from (x_0, x_1) toward the point $(1,1)$. By (2), the point $(1,1)$ lies in the complementary half-space H^C. Hence, if the initial state $(x_0(0), x_1(0))$ lies in H, the distance between $(x_0(t), x_1(t))$ and the line $\hat{L}$ is decreasing in t and reaches 0 in finite time.

Similarly, if $(x_0, x_1) \in H^C$, then $\dot{x}_\theta = \chi - x_\theta(1 - \pi_\theta + \chi)$ for $\theta = 0, 1$. Hence, $(x_0(t), x_1(t))$ converges monotonically toward the point $\left(\frac{\chi}{1+\chi-\pi_0}, \frac{\chi}{1+\chi-\pi_1}\right)$, so long as $(x_0(t), x_1(t))$ remains in H^C. Thus, if the equilibrium is $\left(\frac{\chi}{1+\chi-\pi_0}, \frac{\chi}{1+\chi-\pi_1}\right)$ then the state converges monotonically to the equilibrium starting from any point in H^C, and otherwise the state converges monotonically toward the point $\left(\frac{\chi}{1+\chi-\pi_0}, \frac{\chi}{1+\chi-\pi_1}\right)$ until the state hits the line $\hat{L}$ (which again occurs in finite time).

Next, if $(x_0(t), x_1(t)) \in \hat{L}$ then $\dot{x}_\theta \geq \chi - x_\theta(1 - \pi_\theta + \chi)$ for $\theta = 0, 1$. Hence, if the equilibrium is $\left(\frac{\chi}{1+\chi-\pi_0}, \frac{\chi}{1+\chi-\pi_1}\right)$, then the state converges toward this point from any point in $\hat{L}$. Combining the observations made so far, it follows that when the equilibrium is $\left(\frac{\chi}{1+\chi-\pi_0}, \frac{\chi}{1+\chi-\pi_1}\right)$, it is globally asymptotically stable.

Finally, if $(x_0(t), x_1(t)) \in \hat{L}$ and $(x_0^*, x_1^*) \in \hat{L}$, then the state remains in $\hat{L}$ forever: this follows because, as I have shown, the vector $(\dot{x}_0, \dot{x}_1)$ points toward $\hat{L}$ whenever $(x_0, x_1) \notin \hat{L}$. Next, for any state $(x_0, x_1) \in \hat{L}$, there is a unique mixing probability conditional on observing failure, $s((x_0, x_1))$, such that the vector $(\dot{x}_0, \dot{x}_1)$ is parallel to $\hat{L}$, and in addition the mixing probability $s((x_0, x_1))$ is itself continuous in (x_0, x_1).[30] As the vector $(\dot{x}_0, \dot{x}_1)$ is continuous in (x_0, x_1) and the mixing probability s, it may therefore also be viewed as a continuous function of (x_0, x_1). Furthermore, as any stationary point in $\hat{L}$ is an equilibrium, (x_0^*, x_1^*) is the unique point in $\hat{L}$ such that $(\dot{x}_0, \dot{x}_1) = (0, 0)$. Hence, as $(\dot{x}_0, \dot{x}_1)$ is continuous in (x_0, x_1), it must be that $(\dot{x}_0, \dot{x}_1)$ points toward the steady state, and in addition $(\dot{x}_0(t), \dot{x}_1(t))$ can converge to 0 only if $(x_0(t), x_1(t))$ converges to (x_0^*, x_1^*). Therefore, $(x_0(t), x_1(t))$ must converge to (x_0^*, x_1^*) starting from any initial point in $\hat{L}$. As I have shown that $(x_0(t), x_1(t))$ reaches $\hat{L}$ in finite time starting from any initial point in $[0,1]^2$, it follows that (x_0^*, x_1^*) is globally asymptotically stable. ■

9.11 Proof of Proposition 8

As $(x_0(0), x_1(0))$ is regular, Theorem 1 implies that $(x_0(t), x_1(t))$ is regular for all t. Hence, players experiment with probability 1 after observing a success. On the other hand, a player's

[30] To see this, note that if $s = 0$, then $(\dot{x}_0, \dot{x}_1) = (\chi - x_0(1 - \pi_0 + \chi), \chi - x_1(1 - \pi_1 + \chi))$, which points into H (when $\left(\frac{\chi}{1+\chi-\pi_0}, \frac{\chi}{1+\chi-\pi_1}\right) \in H$, or equivalently when the equilibrium is in $\hat{L}$), and if $s = 1$ then $(\dot{x}_0, \dot{x}_1) = (1 - x_0, 1 - x_1)$, which points into H^C. Denote these vectors by $(\dot{x}_0^0, \dot{x}_1^0)$ and $(\dot{x}_0^1, \dot{x}_1^1)$, and let $\dot{x}_\theta^s = (1-s)\dot{x}_\theta^0 + s\dot{x}_\theta^1$ for $\theta = 0, 1$. By the intermediate value theorem, there is a unique mixing probability $s((x_0, x_1))$ such that $\dot{x}^{s((x_0, x_1))}$ is parallel to L, and $s((x_0, x_1))$ is continuous in (x_0, x_1) because $(\dot{x}_0^0, \dot{x}_1^0)$ and $(\dot{x}_0^1, \dot{x}_1^1)$ are continuous in (x_0, x_1).

posterior after observing a failure at time t is given by

$$p\left(0;x_0\left(t\right),x_1\left(t\right)\right)=\left[1+\frac{1-p}{p}\frac{1}{1-x_1\left(t\right)\pi_1}\right]^{-1}.$$

This posterior is less than p^* at time 0 by (2), and it remains less than p^* until $x_1\left(t\right)$ reaches the value

$$x_1^*=\frac{1}{\pi_1}\left(1-\frac{1-p}{p}\frac{p^*}{1-p^*}\right)<1.$$

(Note that this equation defines the line $\hat{L}$ introduced in the proof of Theorem 7.) Letting T be the first time when $x_1\left(t\right)$ reaches x_1^*, it follows that $\dot{x}_\theta\left(t\right)=-x_\theta\left(t\right)\left(1-\pi_\theta\right)$ for all $t<T$ and $\theta=0,1$. Combined with the initial condition $\left(x_0\left(0\right),x_1\left(0\right)\right)=\left(1,1\right)$, this gives $x_\theta\left(t\right)=\exp\left(-\left(1-\pi_\theta\right)t\right)$ for $\theta=0,1$.

Next, as shown in the proof of Theorem 7, the state $\left(x_0\left(t\right),x_1\left(t\right)\right)$ remains on the line $\hat{L}$ for all $t>T$: that is, $x_1\left(t\right)=x_1^*$ for all $t>T$. It follows that $s\left(t\right)=s$ for all $t>T$, where s is given by $x_1^*=x_1^*\pi_1+\left(1-x_1^*\pi_1\right)s$, or

$$s=\frac{1-\pi_1}{\pi_1}\left(\frac{p}{1-p}\frac{1-p^*}{p^*}-1\right).$$

In addition, for $t>T$, $\dot{x}_0\left(t\right)=s-x_0\left(t\right)$, so $x_0\left(t\right)$ converges monotonically to its steady-state value of s.

Finally, the time T satisfies

$$T=\frac{1}{1-\pi_1}\left[\log\pi_1-\log\left(1-\frac{1-p}{p}\frac{p^*}{1-p^*}\right)\right].$$

Hence,

$$x_0\left(T\right)=\exp\left(-T\right)=\left(\frac{1}{\pi_1}\left(1-\frac{1-p}{p}\frac{p^*}{1-p^*}\right)\right)^{\frac{1}{1-\pi_1}}.$$

In particular, $x_0\left(T\right)<s$ if and only if

$$1-\frac{1-p}{p}\frac{p^*}{1-p^*}<\pi_1\left(\frac{1-\pi_1}{\pi_1}\right)^{1-\pi_1}\left(\frac{p}{1-p}\frac{1-p^*}{p^*}-1\right)^{1-\pi_1}.$$

The right-hand side of this inequality goes to 1 as $\pi_1\to1$, so $x_0\left(T\right)<s$ whenever π_1 is close enough to 1. ■

9.12 Proof of Proposition 9

Fix $\varepsilon \in (0, (\chi - \pi_{\max \Theta^*}) / (1 + \chi - \pi_{\max \Theta^*}))$. Suppose an asymptotically efficient path exists. Then there exists $\bar{K} > 0$ such that if $K > \bar{K}$ then $\left(x_0^K(0), \ldots, x_n^K(0)\right) = (1, \ldots, 1)$ and $\lim_{t \to 0} x_\theta^K(t) < \varepsilon$ (resp., $> 1 - \varepsilon$) for all $\theta \leq \theta^*$ (resp., $> \theta^*$). For such a K, the success rate at $t = 0$ conditional on the event $\theta \leq \theta^*$ equals $(1/a) \sum_{\theta=0}^{\theta^*} p_\theta \pi_\theta$ and the success rate at $t = 0$ conditional on the event $\theta \in \Theta^*$ equals $(1/b) \sum_{\theta \in \Theta^*} p_\theta \pi_\theta$, which is larger. On the other hand, as $t \to \infty$ the success rate conditional on the event $\theta \leq \theta^*$ converges to a number greater than $(1 - \varepsilon) \chi$, while the success rate conditional on the event $\theta \in \Theta^*$ converges to a number less than $\varepsilon + (1 - \varepsilon) \pi_{\max \Theta^*}$, which is smaller. Hence, there must exist a time t^* such that (i) at $t = t^*$, the success rate conditional on the event $\theta \leq \theta^*$ equals the success rate conditional on the event $\theta \in \Theta^*$, and (ii) for all $t > t^*$, the success rate conditional on the event $\theta \leq \theta^*$ is larger than the success rate conditional on the event $\theta \in \Theta^*$.

Now, at $t = t^*$, after observing any sample a player's relative assessment of the probability of the events $\theta \leq \theta^*$ and $\theta \in \Theta^*$ equals the prior probability $a/(a + b)$. Thus, (15) implies that (after observing any sample at $t = t^*$) action 1 is optimal conditional on the event $\theta \in \{1, \ldots, \theta^*\} \cup \Theta^*$. In addition, action 1 is optimal at any state $\theta \notin (\{1, \ldots, \theta^*\} \cup \Theta^*)$. Hence, $\dot{x}_\theta(t^*) = 1 - x_\theta(t^*)$ for all θ. Therefore,

$$\begin{aligned}
\frac{1}{a} \sum_{\theta=0}^{\theta^*} p_\theta \dot{x}_\theta(t^*)(\pi_\theta - \chi) &= \left[\frac{1}{a} \sum_{\theta=0}^{\theta^*} p_\theta (\pi_\theta - \chi)\right] - \left[\frac{1}{a} \sum_{\theta=0}^{\theta^*} p_\theta x_\theta(t^*)(\pi_\theta - \chi)\right] \\
&= \left[\frac{1}{a} \sum_{\theta=0}^{\theta^*} p_\theta (\pi_\theta - \chi)\right] - \left[\frac{1}{b} \sum_{\theta \in \Theta^*} p_\theta x_\theta(t^*)(\pi_\theta - \chi)\right] \\
&< \left[\frac{1}{b} \sum_{\theta \in \Theta^*} p_\theta (\pi_\theta - \chi)\right] - \left[\frac{1}{b} \sum_{\theta \in \Theta^*} p_\theta x_\theta(t^*)(\pi_\theta - \chi)\right] \\
&= \frac{1}{b} \sum_{\theta \in \Theta^*} p_\theta \dot{x}_\theta(t^*)(\pi_\theta - \chi).
\end{aligned}$$

But this implies that, just after time t^*, the success rate conditional on the event $\theta \leq \theta^*$ is smaller than the success rate conditional on the event $\theta \in \Theta^*$, a contradiction. ■

9.13 Proof of Proposition 10

Suppose $\pi_\theta < \chi$ for some efficient state θ, and suppose a simple asymptotically efficient path exists. As in the proof of Proposition 9, for large enough K, at $t = 0$ the success rate in each efficient state is greater than the success rate in each inefficient state, and the situation is reversed for large enough t. Hence, there must exist a time t^* at which the success rates in an inefficient state and an efficient state cross for the first time: that is, a time t^* such that (i) $x_\theta(t^*)(\pi_\theta - \chi) \leq x_{\theta'}(t^*)(\pi_{\theta'} - \chi)$ for all $\theta \leq \theta^* < \theta'$, and (ii) there exists $\varepsilon > 0$ and $\theta \leq \theta^* < \theta'$ such that $x_\theta(t)(\pi_\theta - \chi) > x_{\theta'}(t)(\pi_{\theta'} - \chi)$ for all $t \in (t^*, t^* + \varepsilon)$.

The proof is completed by considering separately the case where $x_\theta(t^*)(\pi_\theta - \chi) = x_{\theta'}(t^*)(\pi_{\theta'} - \chi)$ for all θ, θ' and the case where $x_\theta(t^*)(\pi_\theta - \chi) < x_{\theta'}(t^*)(\pi_{\theta'} - \chi)$ for some θ, θ', and deriving a contradiction in each.

In the first case, the success rate is equal in all states at time t^*, and hence $\dot{x}_\theta(t^*) = 1 - x_\theta(t^*)$ for all θ. But, as in the proofs of Theorem 1 and Proposition 9, this implies that there cannot be a pair of states $\theta < \theta'$ with $x_\theta(t^*)(\pi_\theta - \chi) = x_{\theta'}(t^*)(\pi_{\theta'} - \chi)$ and $x_\theta(t)(\pi_\theta - \chi) > x_{\theta'}(t)(\pi_{\theta'} - \chi)$ for all $t \in (t^*, t^* + \varepsilon)$, a contradiction.

In the second case, there are three states with either (i) $\theta_0 < \theta \leq \theta^* < \theta'$ and

$$x_{\theta_0}(t^*)(\pi_{\theta_0} - \chi) < x_\theta(t^*)(\pi_\theta - \chi) = x_{\theta'}(t^*)(\pi_{\theta'} - \chi)$$

or (ii) $\theta \leq \theta^* < \theta' < \theta_0$ and

$$x_\theta(t^*)(\pi_\theta - \chi) = x_{\theta'}(t^*)(\pi_{\theta'} - \chi) < x_{\theta_0}(t^*)(\pi_{\theta_0} - \chi).$$

Consider the first case (the second is symmetric). Then, as $x_\theta(t)$ is continuous for all θ, for sufficiently small $\varepsilon > 0$,

$$x_{\theta_0}(t^* + \varepsilon)(\pi_{\theta_0} - \chi) < x_{\theta'}(t^* + \varepsilon)(\pi_{\theta'} - \chi) < x_\theta(t^* + \varepsilon)(\pi_\theta - \chi).$$

But then the path is not simple. ■

9.14 Proof of Proposition 11

By assumption, $x_0(0)(\pi_0 - \chi) \leq x_1(0)(\pi_1 - \chi)$. As x_0 and x_1 are continuous, if there exists a time t' with $x_0(t')(\pi_0 - \chi) > x_1(t')(\pi_1 - \chi)$, then there must exist another time t where $x_0(t)(\pi_0 - \chi) = x_1(t)(\pi_1 - \chi)$ but it is not the case that $\dot{x}_0(t)(\pi_0 - \chi) < \dot{x}_1(t)(\pi_1 - \chi)$. As an irregular equilibrium exists, $\pi_1 < \chi$ (by Proposition 14), so $x_0(t)(\pi_0 - \chi) = x_1(t)(\pi_1 - \chi)$ implies $x_0(t) < x_1(t)$. But, by definition of $\dot{x}_\theta(t)$, if $x_0(t)(\pi_0 - \chi) = x_1(t)(\pi_1 - \chi)$ and $x_0(t) < x_1(t)$, then $\dot{x}_0(t) > \dot{x}_1(t)$, and hence $\dot{x}_0(t)(\pi_0 - \chi) < \dot{x}_1(t)(\pi_1 - \chi)$. So there can be no such time t'. ■

9.15 Proof of Proposition 12

It follows immediately from the definition of $\dot{x}_\theta(t)$ and stability from above that $\dot{x}_\theta(t)$ is bounded below 0 for all t such that $x_\theta(t)$ is bounded above x_θ^*. It is also straightforward to argue by contradiction that $x_\theta(t)$ can never cross x_θ^*, completing the proof. ■